Dell™ PowerEdge™ Cluster SE500W Systems

Installation and Troubleshooting Guide

Notes, Notices, and Cautions

NOTE: A NOTE indicates important information that helps you make better use of your computer.

NOTICE: A NOTICE indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.

CAUTION: A CAUTION indicates a potential for property damage, personal injury, or death.

Abbreviations and Acronyms

For a complete list of abbreviations and acronyms, see "Abbreviations and Acronyms."

June 2005

Contents

3 Preparing Your Systems for Clustering

Figures

Tables

1

Introduction

This guide provides information for installing a Dell™ PowerEdge™ Cluster SE500W solution in a corporate enterprise.

The information in this guide includes:

- Basic SCSI cluster installation procedures, which include:
  - Preparing server and storage systems for clustering
  - Cabling the cluster configuration
  - Installing the Microsoft® Windows® 2000 Advanced Server or Windows Server™ 2003 operating system in your cluster configuration

    NOTE: Reference to Windows Server 2003 in this guide implies reference to both Windows Server 2003 Enterprise and Windows Server 2003 Enterprise x64 Editions, unless explicitly stated.
  - Configuring the cluster peripherals, including PERC 4/DC or PERC 4e/DC cards and network adapters
- Installation procedures for installing a SCSI cluster configuration in your corporate network
- Cluster upgrade and maintenance procedures
- Information about MSCS, the clustering software built into the Windows 2000 Advanced Server, and Windows Server 2003 operating systems
- Troubleshooting procedures
- Data form for recording critical cluster configuration information

Intended Audience

This guide was developed for experienced IT professionals who need to install, cable, and configure a PowerEdge Cluster SE500W solution in an enterprise environment and for trained service technicians who perform cluster upgrade and maintenance procedures.

Obtaining More Information

See "Obtaining Technical Assistance" and "Overview" for a general description of PowerEdge clusters and clustering technology.

See "Using MSCS" for an overview of the clustering software built into the Windows 2000 Advanced Server, Windows Server 2003 Enterprise Edition, and Windows Server 2003 x64 Edition operating systems.

Obtaining Technical Assistance

Dell Enterprise Training and Certification is available; see **www.dell.com/training** for more information. This service may not be offered in all locations.

Overview

Clustering uses specific hardware and software to join multiple systems together to function as a single system and provide an automatic failover solution. If one of the clustered systems (also known as *cluster nodes*, or *nodes*) fails, resources running on the failed system are moved (or failed over) to one or more systems in the cluster by the Microsoft® Cluster Service (MSCS) software. MSCS is the failover software component in specific versions of the Windows operating system.

When the failed system is repaired and brought back online, resources automatically transfer back (or *fail back*) to the repaired system or remain on the failover system, depending on how MSCS is configured. See "Failover and Failback" for more information.

NOTE: Reference to Windows Server 2003 in this guide implies reference to both Windows Server 2003 Enterprise and Windows Server 2003 Enterprise x64 Editions, unless explicitly stated.

Virtual Servers and Resource Groups

In a cluster environment, you do not access a physical server; you access a virtual server, which is managed by MSCS. Each virtual server has its own IP address, name, and hard drive(s) in the shared storage system. MSCS manages the virtual server as a *resource group*, which contains the cluster resources. Ownership of virtual servers and resource groups is transparent to users. See "Groups" for more information on resource groups.

When MSCS detects a failed node or failed application, MSCS moves the entire resource group to another node and remaps the virtual server to the new network connection. Users of an application in the virtual server experience only a momentary delay in accessing resources while MSCS re-establishes a network connection to the virtual server and restarts the application.

Quorum Resource

A single disk, which is designated as the quorum resource, maintains the configuration data (including all the changes that have been applied to a cluster database) necessary for recovery when a node fails.

The quorum resource can be any resource with the following attributes:

- Enables a single node to gain and defend its physical control of the quorum resource
- Provides physical storage that is accessible by any node in the cluster
- Uses the Microsoft Windows NT® file system (NTFS)

See "Quorum Disk (Quorum Resource)" and the MSCS online documentation for more information.

NOTE: PowerEdge Cluster SE500W solutions do not support the Majority Node Set (MNS) Quorum resource type.

Shared Storage Systems

Cluster nodes can share access to external storage systems; however, only one of the nodes can own any RAID volume in the external storage system at any time. MSCS controls which node has access to each RAID volume in the shared storage system.

Dell OpenManage™ Array Manager or Dell OpenManage enhanced Storage Manager provides storage management and monitoring for SCSI storage components. See your Array Manager or OMSM documentation for more information.

PowerEdge Cluster SE500W Solution

The PowerEdge Cluster SE500W solution implements two-node clustering technology based on the MSCS software incorporated within the Windows 2000 Advanced Server, and Windows Server 2003 operating systems. This cluster solution provides the following benefits to meet the needs of mission-critical network application programs:

- High availability of system services and resources to network clients
- Redundant storage for application program data
- Failure recovery for cluster-aware applications
- Flexible maintenance capabilities, allowing you to repair, maintain, or upgrade a cluster node without taking the entire cluster offline

Each cluster node is configured with software, storage, and network resources that enable it to monitor and interact with the other node to provide mutually redundant operation. If a cluster node fails for any reason, virtual servers and resource groups are failed over to the healthy cluster node. When the failed node is repaired and brought back online, the virtual servers and resource groups are failed back to the repaired node (if desired).

The cluster nodes, therefore, operate as a single resource, rather than a collection of individual systems. Because the cluster nodes interact in this way, each virtual server appears as a single system to the network clients.

Operating Systems

The PowerEdge Cluster SE500W solution supports two-node cluster configurations that vary in size and performance. Table 1-1 provides an overview of the supported Windows operating systems.

See your operating system documentation for a complete list of features.

Table 1-1. Windows Operating System Features

Windows 2000 Advanced Server	Windows Server 2003 Enterprise Edition	Windows Server 2003 Enterprise x64 Edition
Supports two-node SCSI clusters	Supports two-node SCSI clusters	Supports two-node SCSI clusters
Supports up to 8 GB of RAM per node	Supports up to 32 GB of RAM per node	Supports up to 1 TB RAM per node
Cluster configuration and management using Control Panel utilities	Cluster configuration and management using Configure Your Server (CYS) and Manage Your Server (MYS) wizards	Cluster configuration and management using Configure Your Server (CYS) and Manage Your Server (MYS) wizards
	Metadirectory Services	Metadirectory Services

NOTE: The amount of RAM supported per node also depends on your cluster platform.

NOTE: Running different operating systems in a cluster is supported only during a rolling upgrade. You cannot upgrade to Windows Server 2003 Enterprise x64 Edition. Only a new installation is permitted for Windows Server 2003 Enterprise x64 Edition.

NOTE: MSCS and Network Load Balancing (NLB) features cannot coexist on the same node, but can be used together in a multitiered cluster. For more information, see the Dell PowerEdge Clusters website at www.dell.com/ha or the Microsoft website at www.microsoft.com.

Storage System

Table 1-2 provides a list of supported storage components and the configuration requirements for the cluster nodes and stand-alone systems connected to the storage systems.

Table 1-2. Cluster Storage Requirements

Hardware Components	Minimum Requirement
Supported storage systems	Up to four Dell PowerVault™ 22*x*S storage systems (for the shared disk resource)
Enclosure management modules	Two enclosure management modules (EMMs) per PowerVault 22*x*S enclosure
Power and cooling requirements	Redundant configuration: two fans and dual power supplies per PowerVault 22xS enclosure
Hard drives	At least two SCSI hard drives in each PowerVault 22*x*S enclosure to support hardware-based RAID functionality **NOTE:** The minimum recommendation for an active/passive cluster configuration is two virtual disks. If you are going to use an active/active configuration, at least three virtual disks are recommended.
Cables	Two 1-, 4-, 8-, or 12-m SCSI cables for each PowerVault 22*x*S storage system in the cluster

Hardware and Software Technologies

The PowerEdge Cluster SE500W solution implements the following hardware and software technologies:

- Clustering technology based on the MSCS software in the Windows 2000 Advanced Server, and Windows Server 2003 operating systems
- PERC 4/DC card or PERC 4e/DC card
- Storage management software

Clustering Technology

Clustering is the process of connecting multiple servers together to achieve higher availability and performance. MSCS is a software component in Windows 2000, Advanced Server and an integrated service in Windows Server 2003, which provides failover support for applications and services running on each node.

See "Using MSCS" for more information on the Cluster Service.

NOTE: MSCS and network load balancing (NLB) features cannot coexist on the same cluster node, but can be used together in a multi-tiered cluster configuration. For more information, see the PowerEdge Clusters website located at **www.dell.com/ha** or the Microsoft website located at **www.microsoft.com**.

Cluster Components

- Two PowerEdge systems in a homogeneous pair.

  NOTE: Dell or Microsoft can support only the specific configurations described in the *Dell PowerEdge Cluster SE500W Platform Guide*.

- Up to four PowerVault 22*x*S storage systems, each with dual EMMs and split-bus module.
- Supported network adapters for a public and private LANs.
- Support for RAID 1, 5, 1+0, and 5+0 levels and hot spare drives. If you are using two enclosures, Dell recommends configuring RAID 1 or 1+0 across enclosures for additional redundancy.

  NOTE: RAID 0 and independent drives are possible but are not recommended for a high-availability system because they do not offer data redundancy if a disk failure occurs.

- PERC 4/DC or PERC 4e/DC adapter(s) for the cluster's shared storage.

  NOTE: The PowerEdge Cluster SE500W supports up to two PERC 4/DC or PERC 4e/DC adapters in a single cluster node. Dell does not support use of PERC 4/DC and PERC 4e/DC adapters together in the PowerEdge Cluster SE500W solution.

  NOTE: PERC RAID adapter refers to a PERC 4/DC card or a PERC 4e/DC card.

- Ethernet network adapter or integrated network adapter for node-to-node cluster communications.
- Ethernet crossover or straight cable for node-to-node cluster interconnect.

Figure 1-1 shows a sample configuration of the PowerEdge Cluster SE500W components and their interconnections. See the *Dell PowerEdge Cluster SE500W Platform Guide* for system-specific configuration information.

Figure 1-1. Maximum Configuration of the PowerEdge Cluster SE500W Solution

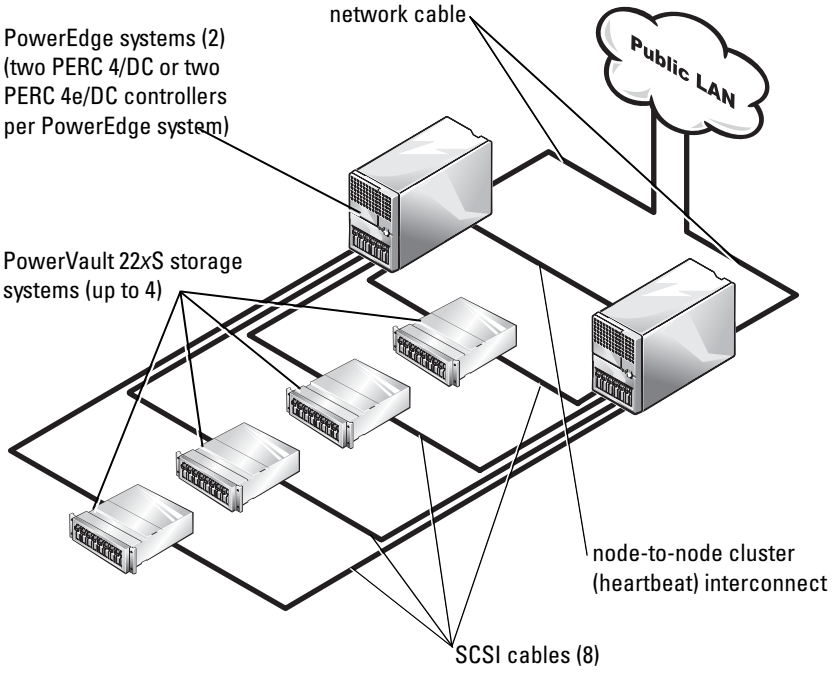

System Requirements

PowerEdge Cluster SE500W configurations require the following hardware and software components:

- Cluster nodes
- Storage
- Interconnects (private network)
- Client network connections (public network)
- Operating system and storage management software

Cluster Nodes

Table 1-3 lists the hardware requirements for the cluster nodes.

Table 1-3. Cluster Node Requirements

Component	Minimum Requirement
Cluster nodes	Two supported PowerEdge systems running Windows 2000 Advanced Server operating system. *or* Two supported PowerEdge systems running Windows Server 2003 operating system.
Processors	At least two processors for each cluster node.
RAM	At least 256 MB of RAM installed on each cluster node for Windows 2000 Advanced Server and Windows Server 2003 Enterprise Edition. At least 512 MB RAM installed on each cluster node for Windows Server 2003 Enterprise x64 Edition.
On-board hard drive controller	For the internal drives, you can use any RAID controller or SCSI adapter supported by the server platform. Two disk drives are required for mirroring (RAID 1) and at least three disk drives are required for disk striping with parity (RAID 5). **NOTE:** Dell strongly recommends that you use hardware-based RAID or software-based disk-fault tolerance for the internal drives. At least 4 GB disk space is required for Windows Server 2003 Enterprise x64 Edition.
Clustered RAID adapter	Up to two PERC 4/DC or two PERC 4e/DC cards for each cluster node. Dell recommends placing your RAID adapters on separate I/O busses to improve availability and performance. See the *Platform Guide* for more information about supported PowerEdge systems.
Network adapters	Minimum of two network interfaces: one for the public network (client LAN connections) and another for the private network (cluster interconnect). The network adapters installed in each cluster node must be identical and supported by the server platform.

Cluster storage

PowerEdge Cluster SE500W configurations support up to four PowerVault 22*x*S storage systems per cluster.

Other Documents You May Need

CAUTION: The *Product Information Guide* provides important safety and regulatory information. Warranty information may be included within this document or as a separate document.

- The *Dell PowerEdge Cluster SE500W Systems Platform Guide* provides information about the systems that support the PowerEdge Cluster SE500W configuration.
- The *Rack Installation Guide* included with your rack solution describes how to install your system into a rack.
- The *Setting Up Your System* document provides an overview of initially setting up your system.
- The *Users Guide* for your PowerEdge or PowerVault system describes system features and technical specifications, SCSI drivers, the System Setup program (if applicable), software support, and the system configuration utility.
- The *Installation and Troubleshooting Guide* for your PowerEdge or PowerVault system describes how to troubleshoot the system and install or replace system components.
- The PERC 4/DC or PERC 4/eDC documentation includes information on the SCSI RAID controller.
- The Dell OpenManage Array Manager or Dell OMSM documentation provides instructions for using the array management software to configure RAID systems.
- Operating system documentation describes how to install (if necessary), configure, and use the operating system software.
- Microsoft Windows 2000 or Windows Server 2003 Cluster Service documentation.
- The *System Administrator's Guide* provides system operation and management operation.
- Documentation for any components you purchased separately provides information to configure and install these options.
- Updates are sometimes included with the system to describe changes to the system, software, and/or documentation.

NOTE: Always read the updates first because they often supersede information in other documents.

- Release notes or readme files may be included to provide last-minute updates to the system documentation or advanced technical reference material intended for experienced users or technicians.

2

Cabling Your Cluster Hardware

Dell™ PowerEdge™ Cluster SE500W configurations require cabling for the storage systems, cluster interconnects, client network connections, and power connections.

Cabling for the Cluster SE500W Solution

The cluster systems and components are interconnected to provide four independent functions as listed in Table 2-1, each of which is described in more detail throughout this section.

Table 2-1. Cluster Cabling Components

Components	Description
Shared storage system	Connects the host-based, RAID controller(s) to the disk enclosure(s). This interconnect is made from the PERC 4/DC or PERC 4e/DC card in each PowerEdge system to the Dell PowerVault™ 22*x*S EMMs, using Dell-qualified U320 SCSI cables.
Cluster interconnect (private network)	Connects the systems to each other to exchange information and status. This connection can be made by using a supported Ethernet network adapter and cabling that is connected to each cluster node. See the *Platform Guide* for a list of supported network adapters for your configuration.
Network connection for public traffic (public network)	Provides a connection between each cluster node and the client network. This connection can be made using an Ethernet network adapter and cabling that is connected to the public network. See the *Platform Guide* for a list of supported network adapters for your configuration.
Power connection	Provides a connection between the power source and the power supplies in your system. By using power strips or Power Distribution Units (PDUs) and separate AC circuits, the cluster can fully utilize the redundant power supplies.

Cabling One PowerVault 22*x*S Shared Storage System to a Cluster SE500W

NOTE: See "Configuring the PowerVault 22xS Storage System for Cluster Mode" for more information about configuring the storage systems.

NOTICE: Do not turn on the systems or the storage system(s) until the split-bus module on the back of the PowerVault system has been set to cluster mode and all cabling is complete.

NOTICE: The asymmetrical, D-shaped SCSI connectors can be damaged if they are forced together when they are not oriented to each other correctly. To avoid connector damage, match the connector shapes (long side-to-long side and short side-to-short side). Insert the cable connector straight (not at an angle) into the card connector.

1. Connect the SCSI connector A (see Figure 2-1) on the back of the PowerVault 22*x*S storage system to the 68-pin connector on the first SCSI cable (see Figure 2-2), and then tighten the retaining screws.

   NOTE: As viewed from the back of the system, the left EMM is the primary EMM and the right EMM is the secondary EMM.

   NOTE: See "Enabling the Cluster Mode Using the PERC RAID Adapter" for more information about cluster-enabling the card.

2. Connect the very high-density connector interface (VHDCI) connector of the first SCSI cable to the channel 0 connector on the cluster-enabled PERC RAID adapter in the first PowerEdge system (see Figure 2-2), and then tighten the retaining screws.

Figure 2-1. PowerVault 22*x*S Back Panel

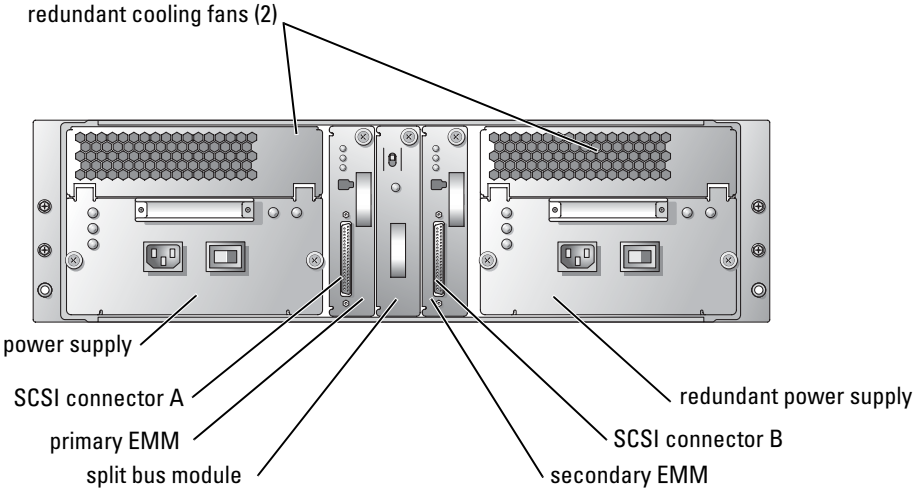

NOTE: Ensure that you securely tighten the retaining screws on all SCSI connectors to ensure a reliable connection.

3 Connect the VHDCI connector of the second SCSI cable (see Figure 2-2) to the channel 0 connector on the cluster-enabled PERC RAID adapter in the second PowerEdge system, and then tighten the retaining screws.

4 Connect the SCSI connector B (see Figure 2-1) on the back of the PowerVault 22*x*S storage system to the 68-pin connector on the second SCSI cable (see Figure 2-2), and tighten the retaining screws.

NOTE: If the PowerVault 22*x*S storage system is disconnected from the cluster, it must be reconnected to the same channel on the same PERC RAID adapter for proper operation.

Figure 2-2. Cabling a Clustered System With One PowerVault 22*x*S Storage System

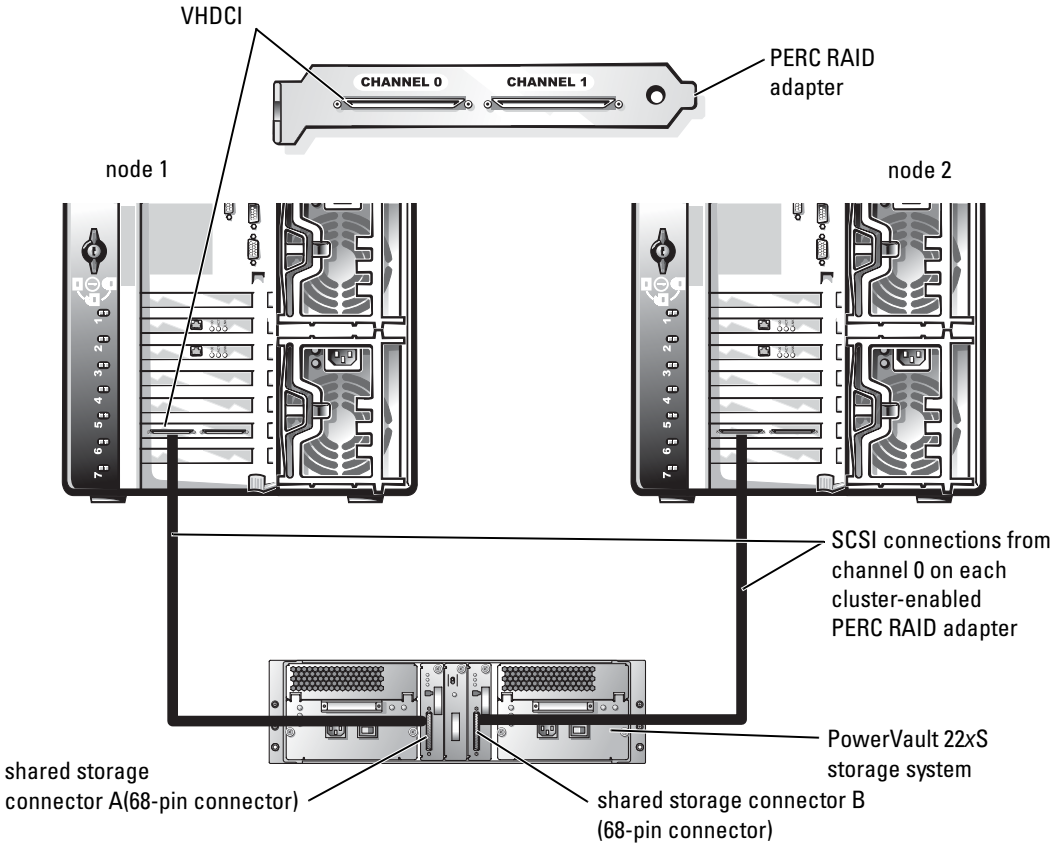

Cabling Two PowerVault 22*x*S Storage Systems to a Cluster SE500W

Connecting the cluster to two PowerVault 22*x*S storage systems is similar to connecting the cluster to a single PowerVault 22*x*S storage system. Connect channel 0 of the PERC RAID adapter in each node to the back of the first storage system, as described in "Cabling One PowerVault 22xS Shared Storage System to a Cluster SE500W." Repeat the process for channel 1 on the controller in each node using a second PowerVault 22*x*S storage system. See Figure 2-3.

With dual storage systems connected to a single PERC RAID adapter, mirroring disk drives from one storage system to another is supported through RAID 1 and 1+0. To protect the cluster applications and your data if an entire storage system fails, Dell strongly recommends using RAID 1 or 1+0.

NOTE: If you have dual cluster-enabled PERC RAID adapters (total of four channels) and only two shared storage systems, you may want to connect one storage system to each controller. If you remove the cable connections, you must reconnect the cables as they were previously connected. To ensure that the cables are reconnected correctly, Dell recommends that you tag or color-code the cables.

Figure 2-3. Cabling Two PowerVault 22*x*S Storage Systems to a PERC RAID adapter

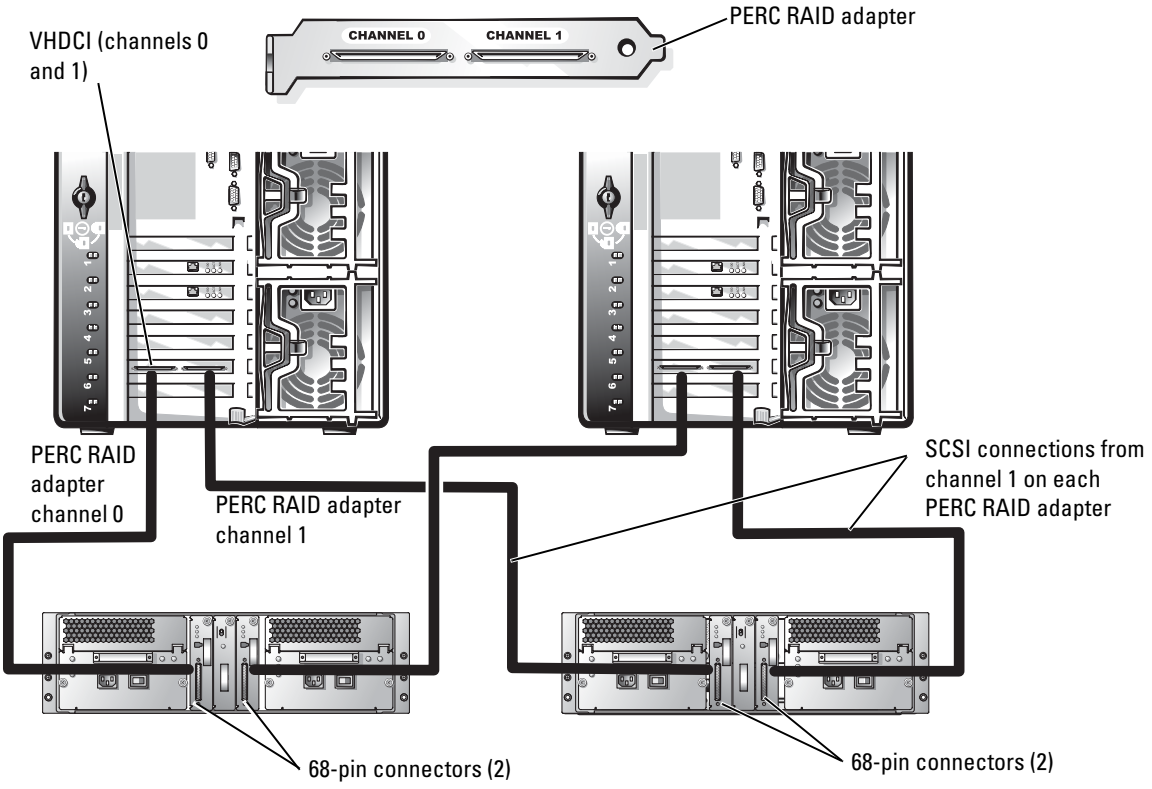

Cabling Three or Four PowerVault 22*x*S Storage Systems to a Cluster SE500W

To connect the cluster to three or four PowerVault 22*x*S storage systems, repeat the process described in the preceding section for a second controller.

NOTICE: If you have dual storage systems that are attached to a second controller, Dell supports disk mirroring between channels on the second controller. However, Dell does not support mirroring disks on one cluster-enabled PERC RAID adapter to disks on another cluster-enabled PERC RAID adapter.

Cabling Your Public and Private Networks

The network adapters in the cluster nodes provide at least two network connections for each node. These connections are described in Table 2-2.

Table 2-2. Network Connections

Network Connection	Description
Public network	All connections to the client LAN.
	At least one public network must be configured for **Mixed mode** for private network failover.
Private network	A dedicated connection for sharing cluster health and status information between the cluster nodes.
	Network adapters connected to the LAN can also provide redundancy at the communications level in case the cluster interconnect fails. See your MSCS documentation for more information on private network redundancy.

Figure 2-4 shows an example of network adapter cabling in which dedicated network adapters in each node are connected to the public network and the remaining network adapters are connected to each other (for the private network).

Figure 2-4. Example of Network Cabling Connection

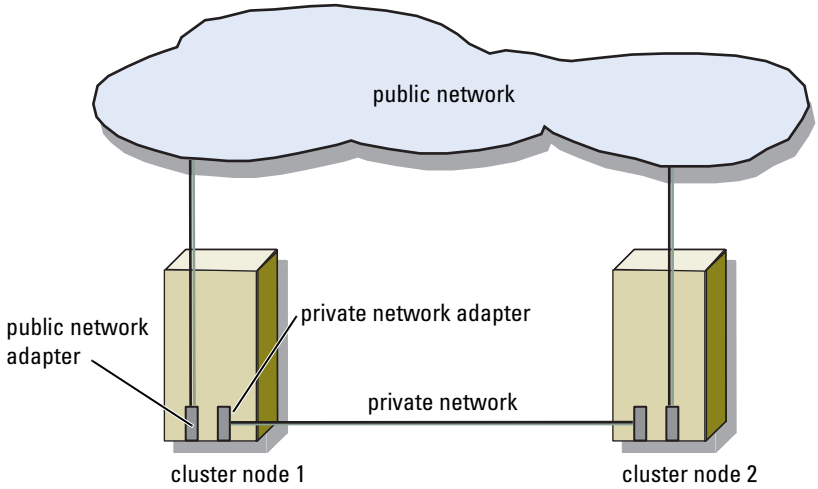

Cabling Your Public Network

The public network connection (client network) to the cluster nodes is provided by a network adapter that is installed in each node. Any network adapter supported by the system running TCP/IP may be used to connect to the public network segments. Additional network adapters may be installed to support additional separate public network segments or to provide redundancy for the public network.

Installing redundant network adapters provides your cluster with a failover connection to the public network. If the primary network adapter or a switch port fails, your cluster will be able to access the public network through the secondary network adapter until the faulty network adapter or switch port is repaired.

Using Dual-Port Network Adapters for Your Private Network

You can configure your cluster to use the public network as a failover for private network communications. However, if dual-port network adapters are used, do not use two ports simultaneously to support both the public and private networks.

Cabling Your Private Network

The private network connection to the cluster nodes is provided by a second or subsequent network adapter that is installed in each node. This network is used for intracluster communications. Table 2-3 lists the required hardware components and connection method for three possible private network configurations.

Table 2-3. Private Network Hardware Components and Connections

Method	Hardware Components	Connection
Network switch	Fast Ethernet or Gigabit Ethernet network adapters and switches	Connect *standard* Ethernet cables from the network adapters in both cluster nodes to a Fast Ethernet or Gigabit Ethernet switch.
Point-to-Point Fast Ethernet	Fast Ethernet network adapters	Connect a *crossover* Ethernet cable between the Fast Ethernet network adapters in both cluster nodes.
Point-to-Point Gigabit Ethernet	Copper Gigabit Ethernet network adapters	Connect a *standard* Ethernet cable between the Gigabit Ethernet network adapters in both cluster nodes.

NOTE: On certain Microsoft® Windows® 2000 Advanced Server, Windows Server™ 2003 configurations, using an Ethernet cable in a point-to-point connection can impact node-to-node communications. See Microsoft Knowledge Base articles 239924, 242430, 254651, and 258750 at **www.microsoft.com** for more information. This issue has been corrected in Windows Server 2003.

NIC Teaming

Network Interface Card (NIC) teaming combines two or more NICs to provide load balancing and/or fault tolerance. Your cluster supports NIC teaming, but only in a public network; NIC teaming is not supported in a private network.

You should use the same brand of NICs in a team, and you cannot mix brands of teaming drivers.

Cabling the Mouse, Keyboard, and Monitor

If you are installing a PowerEdge Cluster SE500W configuration in a Dell rack, your cluster will require a switch box to enable the mouse, keyboard, and monitor for your cluster nodes.

See your rack installation documentation included with your rack for instructions on cabling each cluster node's Keyboard Video Mouse (KVM) to the mouse/keyboard/monitor switch box in the rack.

Power Cabling for the PowerEdge Cluster SE500W Solution

Observe the following cautions when connecting the power cables to the PowerEdge Cluster SE500W.

CAUTION: Although each component of the PowerEdge Cluster SE500W meets leakage current safety requirements, the total leakage current may exceed the maximum that is permitted when the components are used together. To meet safety requirements in the Americas (that is, the United States, Canada, and Latin America), you must use a Type B plug and socket connection for the cluster power to enable the appropriate level of ground protection. In Europe, you must use one or two PDUs or two Type B plug-and- socket connections wired and installed by a qualified electrician in accordance with the local wiring regulations.

CAUTION: Do not attempt to cable the PowerEdge Cluster SE500W to electrical power without first planning the distribution of the cluster's electrical load across available circuits. For operation in the Americas, the PowerEdge Cluster SE500W requires two AC circuits with a minimum capacity of 20 amperes (A) each to handle the electrical load of the system. Do not allow the electrical load of the system to exceed 16 A on either circuit.

CAUTION: For operation in Europe, the PowerEdge Cluster SE500W requires two circuits rated in excess of the combined load of the attached systems. Refer to the ratings marked on the back of each cluster component when determining the total system's electrical load.

See your system and storage system documentation for more information about the specific power requirements for your cluster system's components.

Dell recommends the following guidelines to protect your cluster system from power-related failures:

- For cluster nodes with multiple power supplies, plug each power supply into a separate AC circuit.
- Use uninterruptible power supplies (UPS).

For some environments, you may consider having backup generators and power from separate electrical substations.

Figure 2-5, Figure 2-6, and Figure 2-7 illustrate the proper power cabling of Cluster SE500W components consisting of two PowerEdge systems and two PowerVault 22*x*S storage systems. Each component of the cluster must have power supplied by two or three separate AC circuits—one circuit to each component power supply. Therefore, the primary power supplies of all the PowerEdge Cluster SE500W components are grouped onto one or two circuits and the redundant power supplies are grouped onto a different circuit.

Figure 2-5. Power Cabling Example With Three Power Supplies in the Systems

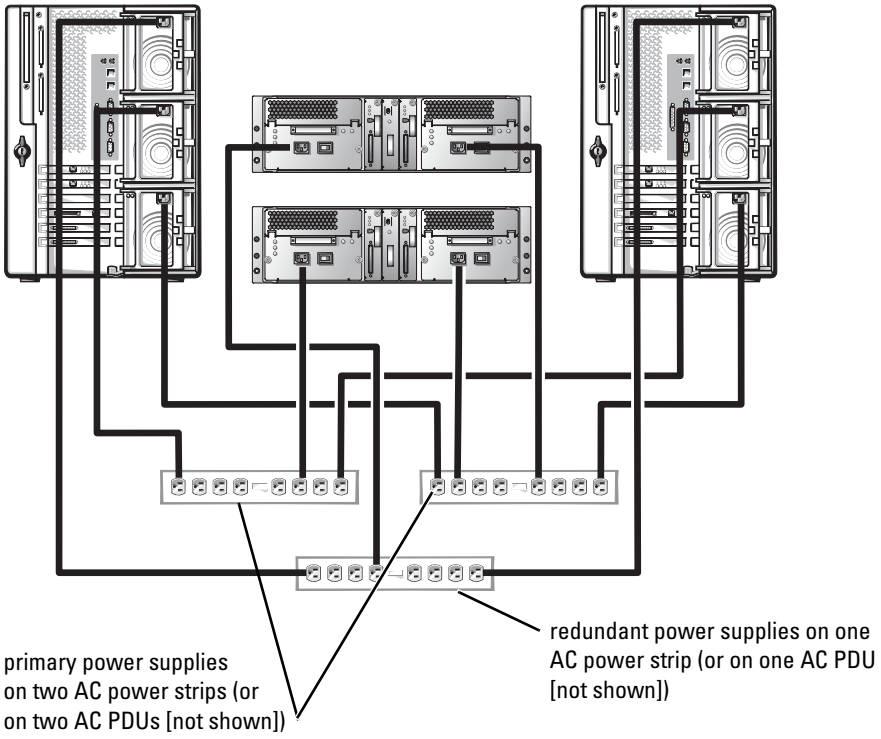

CAUTION: The arrangement of the cluster components in this illustration is intended only to demonstrate the power distribution of the components. Do not stack components as in the configuration shown.

Figure 2-6. Power Cabling Example With One Power Supply in the Systems

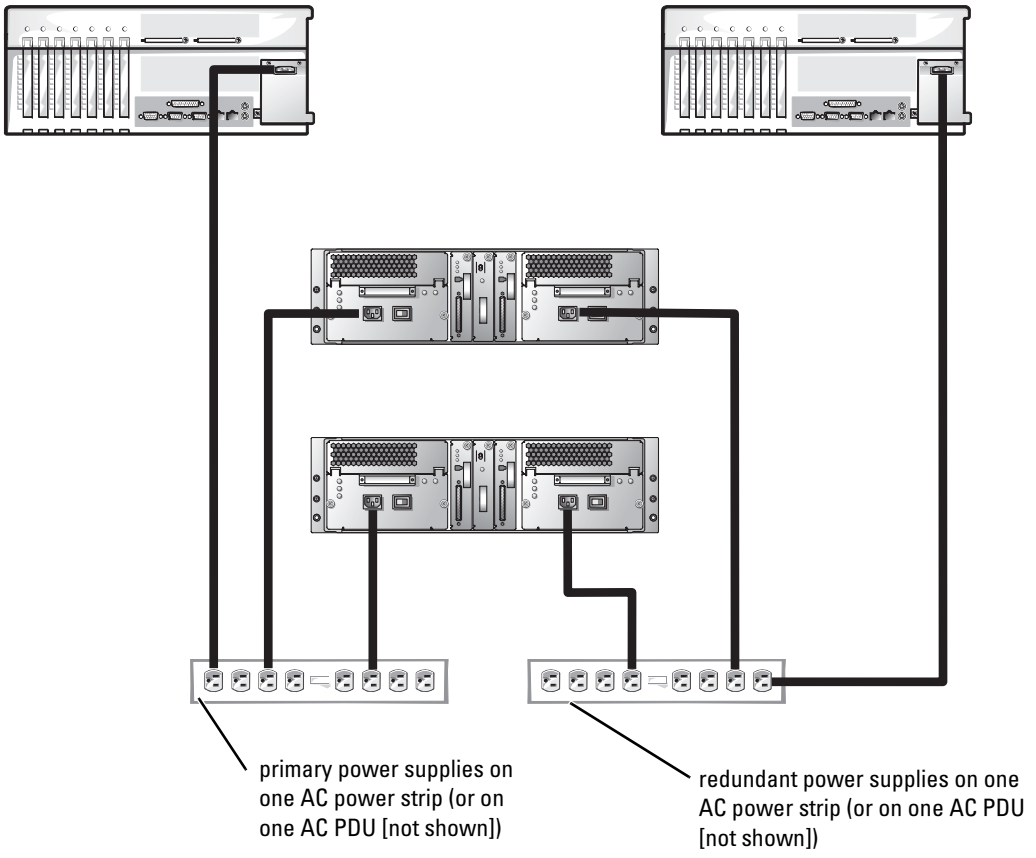

CAUTION: The arrangement of the cluster components in this illustration is intended only to demonstrate the power distribution of the components. Do not stack components as in the configuration shown.

Figure 2-7. Power Cabling Example With Two Power Supplies in the Systems

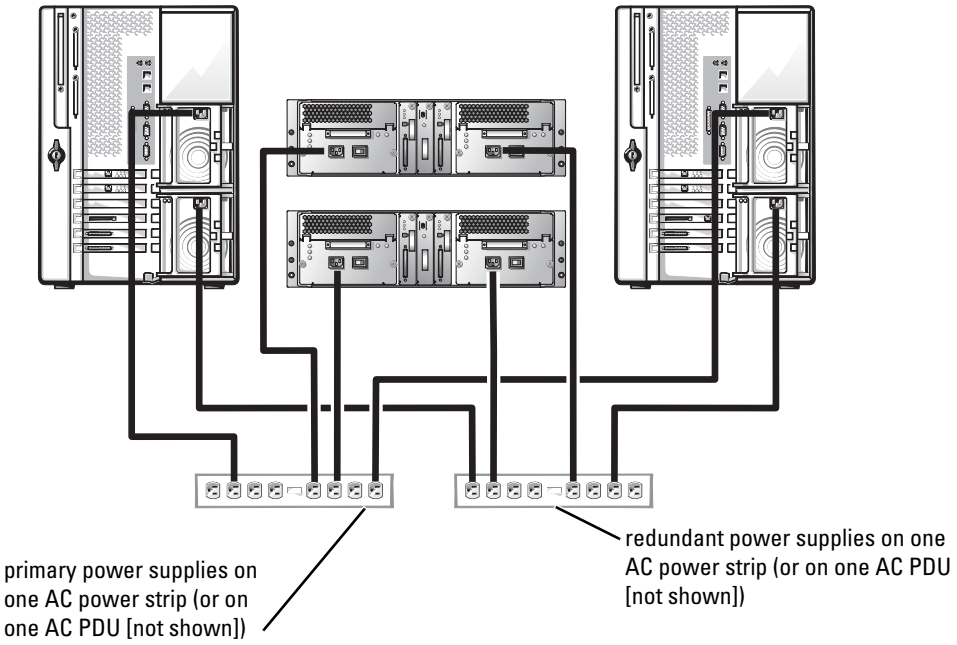

CAUTION: The arrangement of the cluster components in this illustration is intended only to demonstrate the power distribution of the components. Do not stack components as in the configuration shown.

3

Preparing Your Systems for Clustering

This section provides information for preparing your server and storage system hardware for clustering in a Dell™ PowerEdge™ Cluster SE500W configuration.

Before You Begin

1. Ensure that your site can handle the power requirements of the cluster equipment.

   Contact your sales representative for information about your region's power requirements.

CAUTION: Only trained service technicians are authorized to remove and access any of the components inside the system. See your *Product Information Guide* for complete information about safety precautions, working inside the computer, and protecting against electrostatic discharge.

2. Ensure that the following components are installed in each PowerEdge system in the cluster:
   - Network adapters
   - PERC RAID adapters
   - SCSI hard drives
   - Any additional peripheral components
   - RAID controllers for internal drives (optional)
3. Ensure that the following components are installed in each Dell PowerVault™ 22*x*S system in the cluster. See "Installing and Configuring the Shared Storage System."
   - Two enclosure management modules (EMMs)
   - A split-bus module
   - SCSI hard drives
4. Cable the system hardware for clustering.

   See "Cabling Your Cluster Hardware" for more information.
5. Configure the storage system(s) as described in your storage system documentation.
6. Configure the PERC RAID adapters as described in your PERC RAID adapter documentation.
7. If you are using hardware-based RAID for the internal SCSI hard drives, configure the hard drives using the controller's BIOS utility or Dell OpenManage™ Array Manager or Dell OMSM.

Installation Overview

This section provides installation overview procedures for configuring your cluster running the Microsoft® Windows® 2000 Advanced Server, or Windows Server™ 2003 operating system.

1. Ensure that your cluster meets the requirements as described in "Before You Begin."
2. Select a domain model that is appropriate for your corporate network and operating system.

   See "Selecting a Domain Model" for more information.
3. Reserve static IP addresses for your cluster resources and components.

   The resources and components include:

   - Public network
   - Private network
   - Cluster virtual servers

   You will use these IP addresses when you install the Windows operating system and MSCS.

   See "Assigning Static IP Addresses to Your Cluster Resources and Components" for more information.
4. Configure the internal hard drives in your cluster nodes.

   See "Configuring the Internal Drives in Your Cluster Nodes" for more information.
5. Install and configure the Windows operating system.

   The Windows operating system must be installed on both of the cluster nodes. Each cluster node must have its own licensed copy of the Windows operating system and Certificate of Authenticity attached.

   See "Installing and Configuring the Windows Operating System" for more information.
6. Install or update the PERC RAID adapter drivers.

   The PERC RAID adapter drivers allow your cluster nodes to communicate with the shared storage systems.

   See "Updating the Windows Driver for the PERC RAID Adapter" for more information.
7. Install and configure the storage management software.

   See the documentation included with your Array Manager software or available at the Dell Support website (located at **support.dell.com**) for more information.
8. Configure the hard drives on the shared storage system(s).

   See "Configuring and Managing Virtual Disks" for more information.

9 Configure the MSCS software.

The MSCS software is the clustering component of the Windows operating system that provides the failover capabilities for the cluster.

See "Installing and Configuring MSCS" for more information.

10 Verify cluster functionality. Ensure that:

- Your cluster components are communicating properly with each other.
- MSCS is started.

See "Verifying Cluster Functionality" for more information.

11 Verify cluster resource availability.

Use Cluster Administrator to check the running state of each resource group.

See "Verifying Cluster Resource Availability" for more information.

The following sections provide detailed information for each step in the "Installation Overview" that is specific to your Windows operating system.

NOTE: Dell strongly recommends that you use the "PowerEdge Cluster SE500W Solution Data Form" during the installation of your cluster to ensure that all installation steps are completed. The data form is located in "Cluster Data Form."

Selecting a Domain Model

On a cluster running the Windows operating system, both cluster nodes must belong to a common domain or directory model. The following membership configurations are supported:

- Both cluster nodes are member servers in an Active Directory domain.
- Both cluster nodes are domain controllers in an Active Directory domain.
- One cluster node is a domain controller in an Active Directory, and the other cluster node is a member server.

Configuring the Cluster Nodes as Domain Controllers

If a cluster node is not configured as a domain controller and the node cannot contact a domain controller, the cluster node will not be able to authenticate client requests.

If a cluster node is configured as a domain controller, client access to its cluster resources can continue if the cluster node cannot contact other domain controller(s). However, domain controller functions can cause additional overhead, such as log on, authentication, and replication traffic on the cluster nodes.

Configuring the Internal Drives in Your Cluster Nodes

If you have added new hard drives to your system or are setting up the internal drives in a RAID configuration, you must configure the RAID (if applicable) using the RAID controller's BIOS configuration utility or Dell OpenManage Server Assistant before you can install the operating system. For the best balance of fault tolerance and performance, Dell recommends using RAID 1 for the internal hard drives.

RAID 1 is commonly called mirroring because it writes identical data to two separate drives. RAID 1 is the best choice in high-availability applications that require high performance or relatively low data capacity.

See the documentation for your specific RAID controller for more information on RAID configurations.

NOTE: If you are not going to use a hardware-based RAID controller, Dell recommends using the Windows Disk Management tool or Dell OpenManage Array Manager or Dell OMSM to provide software-based redundancy for the Windows system partitions.

Installing and Configuring the Windows Operating System

Before you install the Windows operating system on your cluster nodes, ensure that your cluster meets the requirements listed in "Before You Begin." After you complete these preliminary procedures, install and configure the operating system on your cluster nodes.

The following is an overview sequence for installing the Windows operating system and the cluster management software on the cluster solution.

NOTICE: Windows standby mode and hibernation mode are not supported in cluster configurations. Do not enable either mode.

1. Ensure that your cluster configuration meets the requirements as listed in "Before You Begin."
2. Cable your cluster hardware.

   NOTE: Do *not* connect your cluster nodes to the shared storage systems yet.

   See "Cabling Your Cluster Hardware" for instructions about how to cable your system hardware for cluster configurations, and information about cluster cables, connectors, and components.
3. Install and configure the Windows 2000 Advanced Server or Windows Server 2003 operating system with the latest service pack on each node.

   See the *Platform Guide* for more information about the latest supported service pack.
4. If you are installing Windows Server 2003, go to step 5.

   If you are installing Windows 2000 Advanced Server, select the option to install the Cluster Service files when prompted. You will configure the Cluster Service later.

5. Ensure that the network adapter drivers installed on each cluster node are the latest supported version.
6. Configure the public and private network adapter interconnects in each node, and place the interconnects on separate IP subnetworks using static IP addresses. See "Configuring Windows Networking."

   Public refers to the network adapter that connects the cluster nodes to the client network.

   Private refers to the dedicated network (cluster interconnect) that connects the cluster nodes to each other.
7. Install or update the driver for the PERC RAID adapter in each node and ensure that each PERC card has a different SCSI ID. See "Preparing the PERC RAID Adapter for Clustering" and "Installing the PERC RAID Adapters."
8. Shut down both nodes and connect each node to shared storage. See "Cabling Your Cluster Hardware."
9. Turn on one node and configure shared storage using Dell Storage Management or the PERC RAID adapter BIOS utility. See "Installing and Configuring the Shared Storage System."
10. Reboot node 1.
11. From node 1, write the disk signature and then partition, format, and assign drive letters and volume labels to the hard drives in the storage system using the **Windows Disk Management** application or Array Manager. See "Naming and Formatting Drives on the Shared Storage System."
12. On node 1, verify disk access and functionality on all shared disks.
13. Shut down node 1.
14. Verify disk access by performing the following steps on the other node:
    a. Turn on the node.
    b. Use the PERC RAID adapter BIOS utility to resolve the Non-Volatile Random-Access Memory (NVRAM) mismatch. See your PERC RAID adapter documentation.
    c. Modify the drive letters to match the drive letters on node 1. See "Cluster Data Form."

       This procedure allows the Windows operating system to mount the volumes.
    d. Close and reopen **Disk Management**.
    e. Verify that Windows can see the file systems and the volume labels correctly.
15. Turn on node 1.
16. Install and configure the Cluster Service. See "Installing and Configuring Microsoft Windows 2000 Cluster Service (MSCS)" and "Configuring Microsoft Windows Server 2003 Cluster Service (MSCS)."
17. Install and set up your application programs (optional).
18. Record the cluster configuration using the "Cluster Data Form" (optional).

Configuring Windows Networking

You must configure the public and private networks in each node before you install the Cluster Service. The following sections introduce you to some principles and procedures necessary to the networking prerequisites.

Assigning Static IP Addresses to Your Cluster Resources and Components

A static IP address is an Internet address that a network administrator assigns exclusively to a system or a resource. The address assignment remains in effect until it is changed by the network administrator.

The IP address assignments for the public LAN segments will depend on the configuration of your environment. If the IP assignments are set up correctly, all of the network adapter resources will respond to ping commands and appear online before and after you install MSCS. If the IP assignments are not set up correctly, the cluster nodes may not be able to communicate with the domain. See "Troubleshooting" for more information.

PowerEdge Cluster configurations running the Windows operating system require static IP addresses assigned to hardware and software applications in your cluster, as listed in Table 3-1.

Table 3-1. Applications and Hardware Requiring IP Address Assignments

Application/Hardware	Description
Cluster IP address	The cluster IP address is used for cluster management and must correspond to the cluster name. Because each server has at least two network adapters, the minimum number of static IP addresses required for a cluster configuration is five (one for each network adapter and one for the cluster). Additional static IP addresses are required when MSCS is configured with application programs that require IP addresses, such as file sharing.
Cluster-aware applications running on the cluster	These applications include Microsoft SQL Server, Enterprise Edition Microsoft Exchange Server, and Internet Information Server (IIS). For example, Microsoft SQL Server, Enterprise Edition requires at least one static IP address for the virtual server. (Microsoft SQL Server does not use the cluster's IP address.) Also, each IIS Virtual Root or IIS Server instance configured for failover needs a unique static IP address.

Table 3-1. Applications and Hardware Requiring IP Address Assignments *(continued)*

Application/Hardware	Description
Cluster node network adapters	The network adapters are used to connect to the public and private networks. For cluster operation, two network adapters are required: one network adapter for the public network (LAN/WAN) and another network adapter for the private network (sharing heartbeat information between the cluster nodes). See "Cabling Your Cluster Hardware" for more information about cluster interconnect options. **NOTE:** To ensure cluster operations during a DHCP server failure, Dell recommends using static IP addresses for your cluster. **NOTE:** On certain Windows 2000 Advanced Server configurations, using an Ethernet cable in a point-to-point connection can impact node-to-node communications. See Microsoft Knowledge Base articles 239924, 242430, 254651, and 258750 at **www.microsoft.com** for more information. This issue has been corrected in Windows Server 2003.

Configuring IP Addresses for the Private Network

Dell recommends using the static IP address assignments for the network adapters used for the private network (cluster interconnect). The IP addresses in Table 3-2 are used as examples only.

Table 3-2. Examples of IP Address Assignments

Usage	Cluster Node 1	Cluster Node 2
Public network static IP address (for client and domain controller communications)	192.168.1.101	192.168.1.102
Public network subnet mask	255.255.255.0	255.255.255.0
Private network static IP address cluster interconnect (for node-to-node communications)	10.0.0.1	10.0.0.2
Private network subnet mask	255.255.255.0	255.255.255.0
Default gateway	192.168.1.1	192.168.1.1
Windows Internet Naming Service (WINS) servers	Primary 192.168.1.11	Primary 192.168.1.11
	Secondary 192.168.1.12	Secondary 192.168.1.12
Domain Naming Server (DNS) servers	Primary 192.168.1.21	Primary 192.168.1.21
	Secondary 192.168.1.22	Secondary 192.168.1.22

NOTE: Dell recommends that you do not configure Default Gateway, NetBIOS, WINS, and DNS on your private network. If you are running Windows 2000 Advanced Server or Windows Server 2003 disable NetBIOS on your private network.

If multiple cluster interconnect network adapters are connected to a network switch, ensure that all of the private network's network adapters have a unique address. You can continue the IP address scheme in Table 3-2 with 10.0.0.3, 10.0.0.4, and so on for the private network's network adapters of the other clusters connected to the same switch.

Additional fault tolerance for the LAN segments can be achieved by using network adapters that support adapter teaming or by having multiple LAN segments. To avoid communication problems in the private network, Dell recommends that you do not use dual-port network adapters for the cluster interconnect.

Creating Separate Subnets for the Public and Private Networks

The public and private network's network adapters installed in the same cluster node must reside on separate IP subnetworks. Therefore, the private network used to exchange heartbeat information between the cluster nodes must have a separate IP subnet or a different network ID than the public network, which is used for client connections.

Setting the Network Interface Binding Order for Clusters Running Windows 2000

1. On the Windows 2000 desktop, right-click **My Network Places**, and then click **Properties**.

   The **Network and Dial-up Connections** window appears, displaying all available network interfaces.

2. Click the **Advanced** menu and then click **Advanced Settings**.

   The **Advanced Settings** window appears.

3. In the **Adapters and Bindings** tab in the **Connections** box, ensure that the **Public** connections are designated for **Client access only** or **All communications** are at the top of the list.

   To change the connection order:

   a. Click **Public** or **Private**.

   b. Click the up-arrow or down-arrow to move the connection type to the top or bottom of the **Connections** box.

   c. Click **OK**.

   d. Close the **Network and Dial-up Connections** window.

Setting the Network Interface Binding Order for Clusters Running Windows Server 2003

1. Click the **Start** button, select **Control Panel**, and double-click **Network Connections**.
2. Click the **Advanced** menu, and then click **Advanced Settings**.

   The **Advanced Settings** window appears.

3 In the **Adapters and Bindings** tab, ensure that the **Private** and **Public** connections are at the top of the list.

To change the connection order:

a Click **Public** or **Private**.

b Click the up-arrow or down-arrow to move the connection to the top or bottom of the **Connections** box.

c Click **OK**.

d Close the **Network Connections** window.

Using Dual-Port Network Adapters for the Private Network

You can configure your cluster to use the public network as a failover for private network communications. However, dual-port network adapters are not supported in the private network.

Verifying Cluster Network Communications

To ensure proper cluster operations, the cluster nodes must be able to communicate with each other through the private network (cluster interconnect). This communication involves the exchange of heartbeat messages, whereby the two cluster nodes inquire about each other's status, or "health," and acknowledge each inquiry.

To verify network communications between the cluster nodes:

1 Open a command prompt on each cluster node.

2 At the prompt, type:

```
ipconfig /all
```

3 Press <Enter>.

All known IP addresses for each local server appear on the screen.

4 Issue the **ping** command from each remote system.

Ensure that each local server responds to the **ping** command.

Configuring the Internet Connection Firewall

The Microsoft Windows Server 2003 Enterprise x64 Edition and Windows Server 2003 Enterprise Edition operating systems with Service Pack 1 include an enhanced Internet Connection Firewall that can be configured to block incoming network traffic to a PowerEdge system. To prevent the Internet Connection Firewall from disrupting cluster communications, additional configuration settings are required for PowerEdge systems that are configured as cluster nodes in an MSCS cluster.

Certain network communications are necessary for cluster operations, applications and services hosted by the cluster, and clients accessing those services. If the Internet Connection Firewall is enabled on the cluster nodes, install and run the Security Configuration Wizard, and then configure access for the cluster service and for any applications or services hosted by the cluster and the operating system.

See the following Microsoft Base articles located at the Microsoft Support website at **support.microsoft.com** for more information:

- KB883398 — Internet Connection Firewall
- KB832017 — Network ports used by the Windows Server 2003 operating system

Installing the PERC RAID Adapters

For systems with dual PERC RAID adapters, Dell recommends installing the cards on separate Peripheral Component Interconnect (PCI) buses. Placing the cards on separate buses improves availability and performance.

See the *Platform Guide* for more information about your system's PCI bus configuration.

Updating the Windows Driver for the PERC RAID Adapter

Windows automatically installs a compatible driver for the PERC RAID adapters. The following procedure describes how to update the Windows driver to the PERC family driver.

NOTICE: To ensure proper functionality of the PERC RAID adapter for cluster or internal drive use, this driver *must* be updated. Perform the procedure at the completion of the Windows installation, when the system boots for the first time, and prior to configuring the shared drives.

If a PERC RAID adapter driver CD was not shipped with your system, go to the Dell Support website at **support.dell.com** to download the latest Windows driver for the PERC RAID adapter.

To update the default driver to a PERC RAID adapter driver, perform the following steps:

1. Click the **Start** button, select **Programs**, select **Administrative Tools**, and click **Computer Management**.
2. Select **System Tools**, select **Device Manager**, and click the plus sign (+) to expand **SCSI and RAID controllers**. One or more PERC RAID adapters are listed.
3. Right-click the PERC RAID adapter, select **Properties**, select the **Driver** tab, and then click **Update Driver** to start the Windows Device Driver wizard.
4. Click **Next** to proceed to the **Install Hardware Device Drivers** dialog box.
5. Select **Display a list of known drivers for this device…** and then click **Next**.
6. Click **Have Disk**, insert the diskette or the *Dell OpenManage Server Assistant* CD that contains Dell's updated driver, specify the location of the driver (A:> or D:>), and then click **OK**.
7. Select the appropriate controller (PERC RAID adapter), and then click **Next**.
8. Click **Next** to begin the installation.

9 When the installation is complete, click **Finish** to exit the wizard.

10 Click **Close** to exit the **Properties** window.

11 Click **Yes** to restart the system.

12 Repeat this procedure for cluster node 2.

Installing and Configuring the Shared Storage System

Clustering PowerVault Storage Systems

If you are upgrading an existing PowerVault 22*x*S storage system to meet the cluster requirements for the shared storage system, you may need to install additional hard drives and/or EMMs in the shared storage system. The size and number of drives you add depends on the RAID level you want to use, the number of hard drives installed in your system, and the number of application programs you want to run in your cluster environment. For information about installing hard drives in the PowerVault 22*x*S storage system, see the *Dell PowerVault 220S and 221S System Installation and Troubleshooting Guide*.

NOTE: In cluster mode, the last slot (SCSI ID 15) in the PowerVault 22*x*S is not used; SCSI ID 15 is used for the primary EMM.

Configuring the PowerVault 22*x*S Storage System for Cluster Mode

To ensure that both servers recognize all the drives in the storage system, you must set the split-bus configuration switch to cluster mode on the PowerVault 22*x*S storage system before powering up the storage system.

To configure the PowerVault 22*x*S storage system in cluster mode, perform the following steps:

1 Set the bus configuration switch on the split-bus module to cluster mode (down position). See Figure 3-1.

For more information about the split-bus module, see "Split-Bus Module."

2 Install the split-bus module in the PowerVault 22*x*S storage system.

3 Install the two EMMs in the PowerVault 22*x*S storage system.

See "Enclosure Management Module (EMM)" for basic information about EMMs; see the *Dell PowerVault 220S and 221S Installation and Troubleshooting Guide* for information about installing EMMs.

Split-Bus Module

Your system supports three SCSI bus modes controlled by the split-bus module:

- Joined-bus mode
- Split-bus mode
- Cluster mode

These modes are controlled by the position of the bus configuration switch when the system is turned on. Figure 3-2 illustrates the switch position for each mode.

Figure 3-1. Back-Panel Module Features and Indicators

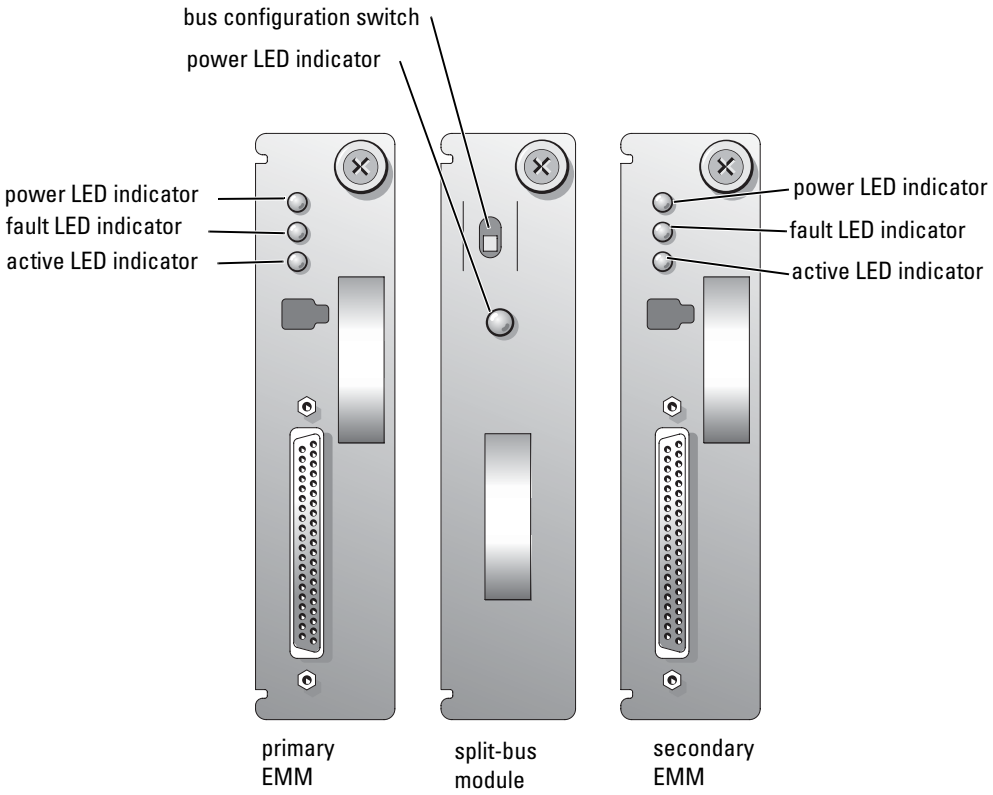

The only difference between cluster mode and joined-bus mode is the SCSI ID occupied by the enclosure services processor. When cluster mode is detected, the processor SCSI ID changes from 6 to 15, allowing a second initiator to occupy SCSI ID 6. As a result, SCSI ID 15 is disabled, leaving 13 available hard drives in cluster mode. For more information about SCSI ID assignments and cluster mode operation, see your *Dell PowerVault 220S and 221S Systems Installation and Troubleshooting Guide*. See Table 3-3 for a description of the split-bus module modes and functions.

NOTE: To change the SCSI bus mode, you must change the position of the bus configuration switch *before* turning on the storage system. Using the bus configuration switch while the system is on does not affect system operation. If you change the bus configuration switch while the system is running, the change will not take effect until you perform the following sequence: shut down the nodes, reboot the storage system, and then turn on the nodes.

Figure 3-2. Bus Configuration Switch Modes

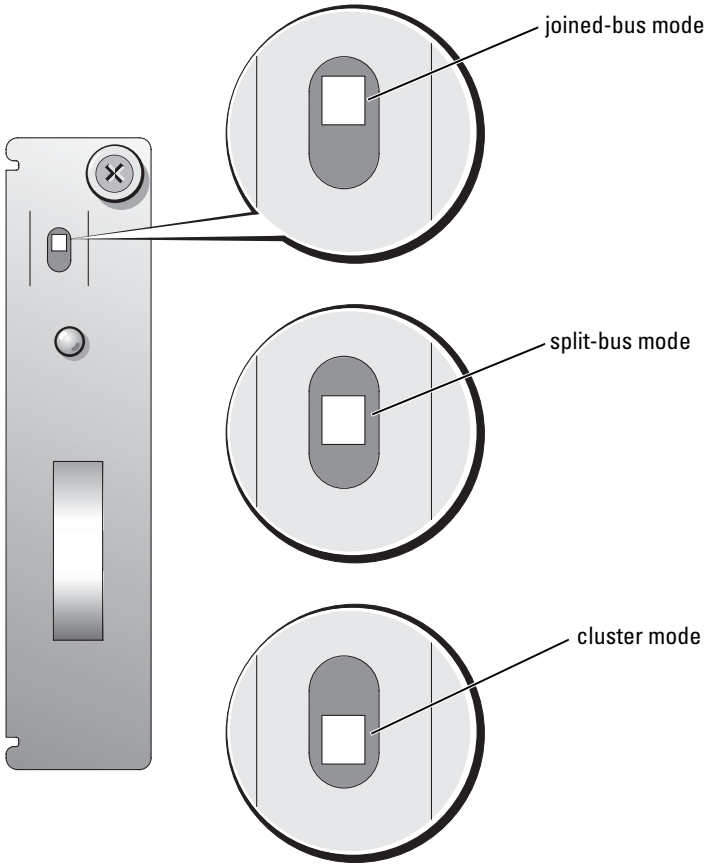

Table 3-3. Split-bus Module Modes and Functions

Mode	Position of Bus Configuration Switch	Function
Joined-bus mode	Up	Low Voltage Differential (LVD) termination on the split-bus module is disabled, electrically joining the two SCSI buses to form one contiguous bus. In this mode, neither the split-bus nor the cluster Light-Emitting Diode (LED) indicators on the front of the enclosure are illuminated.
Split-bus mode	Center	LVD termination on the split-bus module is enabled and the two buses are electrically isolated, resulting in two seven-drive SCSI buses. The split-bus LED indicator on the front of the enclosure is illuminated while the system is in split-bus mode.
Cluster mode	Down	LVD termination is disabled and the buses are electrically joined. The cluster LED on the front of the enclosure is illuminated while the system is in cluster mode.

The split-bus module has only one LED indicator (see Figure 3-1 for location), which is illuminated when the module is receiving power.

Enclosure Management Module (EMM)

The EMM serves two primary functions in your storage system:

- SCSI bus expansion — Acts as a buffer for the SCSI bus, electrically dividing the bus into two independent segments while logically allowing all SCSI bus traffic to pass through it transparently. The buffer improves the quality of the SCSI signals and allows longer cable length connections.
- Management functions — Includes SCSI enclosure services (SES) and SCSI accessed fault-tolerant enclosures (SAF-TE) reporting to the host initiator, control of all enclosure LED indicators, and monitoring of all enclosure environmental elements such as temperature sensors, cooling modules, and power supplies.

A system with redundant enclosure management features two EMMs that are designated as primary and secondary and can be configured in either a cluster, joined-bus, or split-bus mode. A nonredundant configuration consists of one EMM and one SCSI terminator card, and can be configured in a joined-bus mode only. In a redundant system, only one EMM per SCSI bus is active at one time, so only one EMM per SCSI bus can respond to SCSI commands from an initiator.

In joined-bus and cluster modes, if a secondary EMM receives a message that the primary EMM has failed, the fault LED indicator on the primary EMM is illuminated and the condition is reported back to the host initiator. The secondary EMM then becomes active and holds the failed primary in a state of reset until it is replaced. If the primary EMM detects that the secondary has failed, the secondary's fault LED indicator is illuminated and the failed status is reported back to the host initiator.

NOTE: In split-bus mode, each EMM controls half of the enclosure. If one EMM fails in split-bus mode, the second EMM reports the failure but does not assume control of the entire SCSI bus.

The primary EMM is always plugged into the slot on the left (viewed from the back of the system). In a redundant joined-bus configuration, the primary EMM assumes control of all the enclosure functionality. In addition, the primary EMM is the only module that reports the status of the system to the host initiator through SES and SAF-TE protocols. Because the secondary EMM must assume the responsibilities of the primary in the event that the primary fails, both the primary and secondary EMMs are continuously monitoring the status of the system's components.

Preparing the PERC RAID Adapter for Clustering

The warning message shown in Figure 3-3 appears on your screen when you attempt to modify the configuration of the shared storage system on your cluster by using the PERC RAID adapter BIOS configuration utility.

Figure 3-3. Important System Warning

```
 !!!STOP!!!
This operation may change the configuration of disks and
can cause loss of data! Ensure:
1.Peer server is powered up for its controller NVRAM to
be updated. Otherwise, disk configuration should be read
from disk and saved to controller's NVRAM.
2.The second server must not be configuring the disks.
3.There is no I/O activity against shared drives.

Press Any Key To Continue
```

The warning message appears on the screen immediately after activating the PERC RAID adapter BIOS configuration utility by pressing <Ctrl><m> during the system's Power-On Self-Test (POST).

This warning message alerts you to the possibility of data loss if certain precautions are not taken to protect the integrity of the data on your cluster.

NOTICE: To prevent data loss, your cluster must meet the conditions in the following bulleted list before you attempt any data-destructive operation on your shared hard drives.

- Ensure that the peer system is turned on during the operation so that the PERC RAID adapter's NVRAM can be updated with the new configuration information. Alternately, if the peer system is down, you must save the disk configuration to the shared storage system. When you restart the system later, update the peer system's NVRAM from the disk configuration saved to the shared storage system.
- Ensure that the peer cluster node is not currently configuring the shared storage system.
- Ensure that I/O activity does not occur on the shared storage system during the operation.
- Ensure that your PERC RAID adapter firmware is the latest version. See your PERC RAID adapter documentation for information on downloading the latest firmware.

Enabling the Cluster Mode Using the PERC RAID Adapter

Each PERC RAID adapter that is used to connect to a shared storage enclosure must have cluster mode enabled using the PERC RAID adapter's BIOS configuration utility. Enabling cluster mode implements the additional functionality required for the controller to operate in a cluster environment.

NOTICE: If you replace your PERC RAID adapter, ensure that you enable the cluster mode on the replacement PERC RAID adapter and set the SCSI ID to the appropriate value (6 or 7) before you connect the SCSI cables to the shared storage. See "Maintaining Your Cluster" for more information.

For more information about enabling cluster mode, see the PERC RAID adapter documentation, which includes information on the SCSI host adapter.

Setting the SCSI Host Adapter IDs

After cluster mode on the PERC RAID adapter is enabled, the option to change the SCSI ID for both of the adapter's channels will appear. For each shared SCSI bus (a connection from a channel on one system's PERC RAID adapter to the shared storage enclosure to a channel on the second system's PERC RAID adapter), the SCSI IDs for each controller must be different. The default SCSI ID for the PERC RAID adapter is ID 7; thus, the SCSI ID for one of the system's PERC RAID adapters must be changed to ID 6.

For cluster configurations with two PERC RAID adapters in each node connected to shared storage enclosures, set both controllers in one system to SCSI ID 6; that is, one node's pair of PERC RAID adapters utilizes SCSI ID 7 (default) and the other node's pair of PERC RAID adapters is changed to utilize SCSI ID 6.

See the PERC RAID adapter documentation for more information about setting the SCSI host adapter ID number.

NOTICE: If you replace a PERC RAID adapter, you must set the appropriate SCSI ID before you connect the SCSI cables to the shared storage. See "Maintaining Your Cluster" for more information.

Configuring and Managing Virtual Disks

The hard drives in the shared storage system must be set up for clustering. The first step is to configure the RAID levels that you will be using in your cluster. For instructions about setting up a RAID array, see the PERC RAID adapter documentation and the Dell OpenManage Array Manager or Dell OMSM documentation.

All virtual disks, especially if they are used for the quorum resource, should incorporate the appropriate RAID level to ensure high availability. See "Creating the Quorum Resource" for more information on the quorum resource.

NOTE: Dell recommends that you use a RAID level other than RAID 0 (which is commonly called *striping*). RAID 0 configurations provide very high performance, but do not provide the level of availability required for the quorum resource. See the documentation for your storage system for more information about setting up RAID levels for the system.

In a cluster configuration, if multiple NTFS partitions are created on a single virtual disk, these partitions will fail over together. If you plan to run cluster-aware applications on each cluster node, you must create at least two separate virtual disks to ensure that the applications can fail over independently.

"Naming and Formatting Drives on the Shared Storage System" describes how to assign drive letters to the shared hard drives in a cluster installation. For instructions on partitioning and formatting the shared storage system's hard drives, see the appropriate operating system documentation and the PERC RAID adapter documentation.

Windows 2000 and Windows Server 2003 Dynamic Disks and Volumes

The Windows operating system does not support dynamic disks (upgraded disks) or volumes as shared cluster storage. If the shared cluster storage is configured as a dynamic disk, the Cluster Configuration wizard is not able to discover the disks, preventing the cluster and network clients from accessing the disks.

Naming and Formatting Drives on the Shared Storage System

After the virtual disks are created, write the disk signature, assign drive letters to the virtual disks, and then format the drives as NTFS drives. Format the drives and assign drive letters from only one cluster node.

NOTICE: Accessing the hard drives from multiple cluster nodes may corrupt the file system.

Assigning Drive Letters and Mount Points

A mount point is a drive attached to an empty folder on an NTFS volume. A mount point functions the same as a normal drive but is assigned a label or name instead of a drive letter. Using mount points, a cluster can support more shared disks than the number of available drive letters.

The cluster installation procedure does not automatically add the mount point into the disks managed by the cluster. To add the mount point to the cluster, create a physical disk resource in the cluster resource group for each mount point. Ensure that the new physical disk resource is in the same cluster resource group and is dependent on the root disk.

NOTE: Mount points are only supported in MSCS on the Windows Server 2003 operating systems. When mounting a drive to an NTFS volume, do not create mount points from the quorum resource or between the clustered disks and the local disks. Mount points must be in the same cluster resource group and must be dependent on the root disk.

NOTICE: If the disk letters are manually assigned from the second node, the shared disks are simultaneously accessible from both nodes. To ensure file system integrity and prevent possible data loss before you install the MSCS software, prevent any I/O activity to the shared drives by performing the following procedure on one node at a time and ensuring that the other node is shut down.

Before installing MSCS, ensure that both nodes have the same view of the shared storage systems. Because each node has access to hard drives that are in a common storage array, each node must have identical drive letters assigned to each hard drive. Up to 22 logical drive letters (E through Z) can be used for the shared storage systems.

NOTE: Drive letters A through D are reserved for the local system.

The number of drive letters required by individual servers in a cluster may vary.
Dell recommends that the shared drives be named in reverse alphabetical order beginning with the letter *z*.

To assign drive letters and format drives on the shared storage system, perform the following steps:

1. With node 2 shut down, open **Disk Management** on node 1.

   NOTE: For Windows 2000 Advanced Server, Array Manager launches instead of Disk Management if Array Manager is installed on your system.

2. Allow Windows to enter a signature on all new physical or logical drives.

   NOTE: Do not create dynamic disks on your hard drives.

3. Locate the icon for the first unnamed, unformatted drive on the shared storage system.
4. Right-click the icon and select **Create** from the submenu.

   If the unformatted drives are not visible, verify the following:

   - The latest version of the PERC RAID adapter driver is installed.
   - The storage system is properly cabled to the servers.
   - The split-bus module on the PowerVault 22*x*S is set to cluster mode.

5 In the dialog box, create a partition the size of the entire drive (the default) and then click **OK**.

NOTE: The MSCS software allows only one node to access a logical drive at a time. If a logical drive is partitioned into multiple disks, only one node is able to access all the partitions for that logical drive. If each node must access a separate disk, two or more logical drives must be present in the storage system.

6 Click **Yes** to confirm the partition.

7 With the mouse pointer on the same icon, right-click and select **Change Drive Letter and Path** from the submenu.

8 Assign a drive letter to an NTFS volume or create a mount point.

To assign a drive letter to an NTFS volume:

a Click **Edit** and select the letter you want to assign to the drive (for example, *z*).

b Click **OK**.

c Go to step 9.

To create a mount point:

a Click **Add**.

b Click **Mount in the following empty NTFS folder**.

c Type the path to an empty folder on an NTFS volume, or click **Browse** to locate it.

d Click **OK**.

e Go to step 9.

9 Click **Yes** to confirm the changes.

10 Right-click the drive icon again and select **Format** from the submenu.

11 Under **Volume Label**, enter a descriptive name for the new volume; for example, `Disk_Z` or `Email_Data`.

12 In the dialog box, change the file system to **NTFS**, select **Quick Format**, and click the **Start** button.

NOTE: The NTFS file system format is required for shared-disk resources under MSCS.

13 Click **OK** at the warning.

14 Click **OK** to acknowledge that the format is complete.

15 Click **Close** to close the dialog box.

16 Repeat step 3 through step 15 for each remaining drive.

17 Close **Disk Management**.

18 Shut down node 1.

19 Turn on node 2.

20 On node 2, open **Disk Management**.

21 Ensure that the drive letters for node 2 are correct.

To modify the drive letters on node 2, repeat step 7 through step 9.

Installing and Configuring MSCS

MSCS is a software component in Windows 2000 and an integrated service in Windows Server 2003. MSCS performs the basic cluster functionality, which includes membership, communication, and failover management. When MSCS is installed properly, the service starts on each node and responds automatically if one of the nodes fails or goes offline. To provide application failover for the cluster, the MSCS software must be installed on both cluster nodes. See "Using MSCS" for more information.

NOTE: For systems with split backplane modules installed, the cluster installation tries to use the logical drives on the secondary backplane as cluster disks. Because these drives are not accessible to all nodes in the cluster, ensure that they are removed from the cluster after the installation is complete.

NOTE: In Windows Server 2003, mapping a network drive to the same drive letter as a cluster disk resource renders the cluster disk inaccessible from Windows Explorer on the host. Ensure that mapped network drives and cluster disks are never assigned the same drive letter.

Verifying Cluster Readiness

To ensure that your server and storage systems are ready for MSCS installation, ensure that these systems are functioning correctly and verify the following:

- All cluster servers are able to log on to the same domain.
- The shared disks are partitioned and formatted, and the same drive letters that reference logical drives on the shared storage system are used on each node.
- For each attached PowerVault 22*x*S storage system, the split-bus module is set to cluster mode before power-up.
- Cluster mode is enabled on all PERC RAID adapters connected to shared storage.
- The controller's SCSI IDs (6 or 7) on each node are different.
- All peer PERC RAID adapters are connected to the same PowerVault system through the same channel number.
- All IP addresses and network names for each system node are communicating with each other and the rest of the network. The private IP addresses should not be accessible from the LAN.

Creating the Quorum Resource

When you install Windows 2000 Advanced Server in your cluster, the software installation wizard prompts you for a location to install the quorum resource. When you install Windows Server 2003 in your cluster, the software installation wizard automatically selects an NTFS disk as the quorum disk, which you can modify later using **Cluster Administrator**. To prevent quorum disk corruption, Dell and Microsoft recommend that you do not place applications or data on the disk.

Creating a Partition for the Quorum Disk Resource

Dell recommends creating a separate partition—approximately 1 GB in size—for the quorum resource.

When you create the partition for the quorum resource:

- Format the partition with NTFS.
- Use the partition exclusively for your quorum logs.
- Do not store any application data or user data on the quorum resource partition.
- To properly identify the quorum resource, Dell recommends that you assign the drive letter Q to the quorum resource partition.

Dell does not recommend using the remainder of the virtual disk for other cluster resources. If you do use the space for cluster resources, be aware that when you create two volumes (partitions) on a single virtual disk, they will both fail over together if a server fails.

Preventing Quorum Disk Failure

Because the quorum disk plays a crucial role in cluster operation, losing a quorum disk causes the entire cluster to fail. To prevent cluster failure, configure the quorum disk on a RAID volume in the shared storage system.

NOTICE: Dell recommends that you use a RAID level other than RAID 0, which is commonly called striping. RAID 0 configurations provide very high performance, but they do not provide the level of availability required for the quorum resource.

Configuring Cluster Networks Running Windows 2000

When you install and configure a cluster running Windows 2000 Advanced Server, the software installation wizard prompts you to identify the public and private network segments connected to your cluster nodes. Dell recommends the following configuration, which provides added fault tolerance for the private network:

1. Set the private network (cluster interconnect) to **Use for internal communications only**.
2. Name this network **Private**.

3 Set the client system's public network segment(s) to **All communications**.

This setting provides a redundant path for the cluster-to-cluster communication in the event the private network fails.

4 For all subsequent network adapters, set each network adapter to **Client use only** or **All communications**.

5 Set the priority of the networks so that the network you designated as **Private** has the highest priority for internal communications.

You can set the priority of the networks when you install MSCS or when using Cluster Administrator software.

Configuring Cluster Networks Running Windows Server 2003

When you install and configure a cluster running Windows Server 2003, the software installation wizard automatically assigns and configures the public and private networks for your cluster. You can rename a network, allow or disallow the cluster to use a particular network, or modify the network role using **Cluster Administrator**. Dell recommends that you configure at least one network for the cluster interconnect (private network) and one network for all communications.

Installing and Configuring Microsoft Windows 2000 Cluster Service (MSCS)

After the Windows 2000 setup files are copied to the cluster node, Windows 2000 reboots and enters the second phase of installation.

During this phase, the Windows 2000 setup prompts you for the product key, licensing mode, system name, and additional information before starting the Windows 2000 Components program. From this dialog box, you can select the Windows 2000 core components for installation. Selecting **Cluster Service** copies the required files to the system.

If your operating system was preinstalled by Dell or you did not select **Cluster Service** when you installed the operating system, you can install the Cluster Service later by running the **Add/Remove Components** from the **Control Panel**.

After the Windows 2000 setup completes and all clustering prerequisites are met, run the Cluster Configuration wizard to complete the installation, add the other node to the cluster, and configure cluster resources.

To access the Cluster Configuration wizard:

1 From node 1, click the **Start** button and select **Run**.

2 In the **Run** box, type the following:

```
C:\windows\cluster\cluscfg.exe
```

3 Click **OK**.

4 Follow the directions in the **Cluster Configuration Wizard** window.

5 Add the remaining node to the cluster:

a From node 2, click the **Start** button and select **Run**.

b In the **Run** box, type `C:\windows\cluster\cluscfg.exe`.

c Click **OK**.

d When prompted, select **The second or next node in the cluster**.

You are joining an existing cluster.

Configuring Microsoft Windows Server 2003 Cluster Service (MSCS)

The cluster setup files are automatically installed on the system disk. To create a new cluster:

1 From either node, click the **Start** button, select **Programs→ Administrative Tools,** and then double-click **Cluster Administrator**.

2 From the **File** menu, select **Open Connection**.

3 In the **Action** box of the **Open Connection to Cluster**, select **Create new cluster**.

The **New Server Cluster Wizard** appears.

4 Click **Next** to continue.

5 Follow the procedures in the wizard, and then click **Finish**.

6 Add the second node to the cluster.

a Turn on the remaining node.

b Click the **Start** button, select **Programs→ Administrative Tools,** and double-click **Cluster Administrator**.

c From the **File** menu, select **Open Connection**.

d In the **Action** box of the **Open Connection to Cluster**, select **Add nodes to cluster**.

e In the **Cluster or server name** box, type the name of the cluster or click **Browse** to select an available cluster from the list, and then click **OK**.

The **Add Nodes Wizard** window appears.

- If the Add Nodes Wizard *does not* generate a cluster feasibility error, go to step f.
- If the Add Nodes Wizard generates a cluster feasibility error, go to "Adding Cluster Nodes Using the Advanced Configuration Option."

f Click **Next** to continue.

g Follow the procedures in the wizard, and then click **Finish**.

Adding Cluster Nodes Using the Advanced Configuration Option

If you are adding additional nodes to the cluster using the Add Nodes Wizard and the nodes are not configured with identical internal storage devices, the wizard may generate one or more errors while checking cluster feasibility in the **Analyzing Configuration** menu. If this situation occurs, select **Advanced Configuration Option** in the Add Nodes Wizard to add the nodes to the cluster.

To add the nodes using the **Advanced Configuration Option**:

1. From the **File** menu in Cluster Administrator, select **Open Connection**.
2. In the **Action** box of the **Open Connection to Cluster**, select **Add nodes to cluster**, and click **OK**.

   The **Add Nodes Wizard** window appears.
3. Click **Next**.
4. In the **Select Computers** menu, click **Browse**.
5. In the **Enter the object names to select (examples)**, type the names of one to seven systems to add to the cluster, with each system name separated by a semicolon.
6. Click **Check Names**.

   The Add Nodes Wizard verifies and underlines each valid system name.
7. Click **OK**.
8. In the **Select Computer**s menu, click **Add**.
9. In the **Advanced Configuration Options** window, click **Advanced (minimum) configuration**, and then click **OK**.
10. In the **Add Nodes** window, click **Next**.
11. In the **Analyzing Configuration** menu, Cluster Administrator analyzes the cluster configuration.

    If Cluster Administrator discovers a problem with the cluster configuration, a warning icon appears in **Checking cluster feasibility**. Click the plus sign (+) to review any warnings, if needed.
12. Click **Next** to continue.
13. In the **Password** field of the **Cluster Service Account** menu, type the password for the account used to run the Cluster Service, and click **Next**.

    The **Proposed Cluster Configuration** menu appears with a summary with the configuration settings for your cluster.

14 Click **Next** to continue.

The new systems (hosts) are added to the cluster. When completed, **Tasks completed** appears in the **Adding Nodes to the Cluster** menu.

NOTE: This process may take several minutes to complete.

15 Click **Next** to continue.

16 In the **Completing the Add Nodes Wizard**, click **Finish**.

Verifying Cluster Functionality

To verify cluster functionality, monitor the cluster network communications to ensure that your cluster components are communicating properly with each other. Also, verify that MSCS is running on the cluster nodes.

Verifying MSCS Operation

After you install MSCS, verify that the service is operating properly.

1 Click the **Start** button and select **Programs**→ **Administrative Tools**, and then select **Services**.

2 In the **Services** window, verify the following:

- In the **Name** column, **Cluster Service** appears.
- In the **Status** column, Cluster Service is set to **Started**.
- In the **Startup Type** column, Cluster Service is set to **Automatic**.

Verifying Cluster Resource Availability

In the context of clustering, a resource is a basic unit of failover management. Application programs are made up of resources that are grouped together for recovery purposes. All recovery groups, and therefore the resources that comprise the recovery groups, must be online (or in a ready state) for the cluster to function properly.

To verify that the cluster resources are online:

1 Start **Cluster Administrator** on the monitoring node.

2 Click the **Start** button and select **Programs**→ **Administrative Tools (Common)**→ **Cluster Administrator**.

3 Open a connection to the cluster and observe the running state of each resource group. If a group has failed, one or more of its resources might be offline.

Installing Applications in the Cluster Group

The Cluster Group contains a network name and IP address resource, which is used to manage the cluster. Because the Cluster Group is dedicated to cluster management and for best cluster performance, Dell recommends that you do not install applications in this group.

Troubleshooting Failed Resources

Troubleshooting the failed resources is beyond the scope of this document, but examining the properties of each resource and ensuring that the specified parameters are correct are the first two steps in this process. In general, if a resource is offline, it can be brought online by right-clicking the resource and selecting **Bring Online** from the pull-down menu. See the documentation and online help for Windows 2000 Advanced Server or Windows Server 2003 for information about troubleshooting resource failures.

Obtaining More Information

See Microsoft's online help for configuring the Cluster Service.

See "Using MSCS" for more information about the Cluster Service.

4

Installing Your Cluster Management Software

This section provides information on configuring and administering your cluster using Microsoft® Cluster Administrator. Microsoft provides Cluster Administrator as a built-in tool for cluster management.

Microsoft Cluster Administrator

Cluster Administrator is Microsoft's tool for configuring and administering a cluster. The following procedures describe how to run Cluster Administrator locally on a cluster node and how to install the tool on a remote console.

Launching Cluster Administrator on a Cluster Node

1. Click the **Start** button and select **Programs**.
2. Select **Administrative Tools**.
3. Select **Cluster Administrator**.

Running Cluster Administrator on a Remote Console

You can administer and monitor the Cluster Service remotely by installing the Windows Administration Tools package and Cluster Administrator on a remote console (or management station) running the Microsoft Windows® operating system. Cluster Administrator is part of the Administration Tools package, which is included with the following operating systems:

- Windows 2000 Advanced Server
- Windows Server™ 2003 Enterprise Edition
- Windows Server 2003 Enterprise x64 Edition

The Windows 2000 Administration Tools can only be installed on systems running Windows 2000. Additionally, the Windows 2003 Administrative Tools can only be installed on systems running Windows XP (with Service Pack 1 or later) and Windows Server 2003.

To install Cluster Administrator and the Windows Administration Tools package on a remote console, perform the following steps:

1. Select a system that you wish to configure as the remote console.
2. Identify the operating system that is currently running on the selected system.
3. Insert the appropriate operating system CD into the system's CD drive:
   - *Microsoft Windows 2000 Advanced Server CD*
   - *Windows Server 2003 Enterprise Edition CD*
   - *Windows Server 2003 Enterprise x64 Edition CD*
4. Open an Explorer window, navigate to the system's CD drive and double-click the \i386 directory.
5. If you inserted the *Windows 2000 Advanced Server CD* or *Windows Server 2003 Enterprise Edition CD*, double-click **ADMINPAK.MSI**.

   If you inserted the *Windows Server 2003 Enterprise x64 Edition CD*, double-click **WADMINPAK.MSI**.
6. Follow the instructions on your screen to complete the installation.

Launching Cluster Administrator on a Remote Console

Perform the following steps on the remote console:

1. Ensure that the Windows Administrative Tools package was installed on the system.
2. Click the **Start** button and select **Programs**.
3. Select **Administrative Tools**.
4. Select **Cluster Administrator**.

Installing Cluster Administrator for Windows Clusters on a Remote Console

You cannot install the Windows 2000 or Windows Server 2003 Administration Tools package on clients running any version of Windows NT® 4.0. However, a Windows 2000 Advanced Server or Windows Server 2003, Enterprise Edition cluster can be remotely administered using the Cluster Administrator included with Windows NT 4.0 operating systems with limited support.

See your Windows NT 4.0 operating system documentation for more information about the installation of Cluster Administrator on a remote client.

Cluster Administration and Monitoring

When using Cluster Administrator provided with Windows NT 4.0 operating systems on a system running Windows NT 4.0, Cluster Administrator may generate error messages if the software detects Windows 2000 or Windows Server 2003 cluster resources. Dell strongly recommends using client systems running Windows 2000 or Windows Server 2003 with the appropriate Administrator Pack for cluster administration and monitoring.

5

Using MSCS

This section provides information about Microsoft® Cluster Service (MSCS). This section is intended to be an overview of MSCS and provides information about the following:

- Cluster objects
- Cluster networks
- Network interfaces
- Cluster nodes
- Groups
- Cluster resources
- Failover and failback

For information about specific MSCS procedures, see the MSCS online help.

NOTE: In this guide and in other cluster documentation, the quorum resource is also referred to as the quorum disk.

Cluster Objects

Cluster objects are the physical and logical units managed by MSCS. Each object is associated with the following:

- One or more properties, or attributes, that define the object and its behavior within the cluster
- A set of cluster control codes used to manipulate the object's properties
- A set of object management functions used to manage the object through MSCS

Cluster Networks

A network performs one of the following roles in a cluster:

- A network that carries internal cluster communication
- A public network that provides client systems with access to cluster application services
- A public-and-private network that carries both internal cluster communication and connects client systems to cluster application services
- Neither a public nor private network that carries traffic unrelated to cluster operation

Preventing Network Failure

The Cluster Service uses all available private and public-and-private networks for internal communication. Configure multiple networks as private or public-and-private to protect the cluster from a single network failure. If there is only one such network available and it fails, the cluster nodes stop communicating with each other. When two nodes are unable to communicate, they are partitioned, and the Cluster Service automatically shuts down on one node. While this shutdown guarantees the consistency of application data and the cluster configuration, it can make cluster resources unavailable.

For example, if each node has only one network adapter, and the network cable on one of the nodes fails, each node (because it is unable to communicate with the other) attempts to take control of the quorum disk. There is no guarantee that the node with a functioning network connection will gain control of the quorum disk. If the node with the failed network cable gains control, the entire cluster is unavailable to network clients. To avoid this problem, ensure that all nodes have at least two networks and are configured to use both networks for the private network (internal communications).

Node-to-Node Communication

The Cluster Service does not use public-only networks for internal communication. For example, a cluster has Network A configured as private and Network B configured as public. If Network A fails, the Cluster Service does not use Network B because it is public; the nodes stop communicating, and one node terminates its Cluster Service.

Network Interfaces

The Microsoft® Windows® operating system keeps track of all network adapters in a server cluster. This tracking system allows you to view the state of all cluster network interfaces from a cluster management application, such as Cluster Administrator.

Cluster Nodes

A cluster node is a system in a server cluster that has a working installation of the Windows operating system and the Cluster Service.

Cluster nodes have the following characteristics:

- Every node is attached to one or more cluster storage devices. Each cluster storage device attaches to one or more disks. The disks store all of the cluster's configuration and resource data. Each disk can be owned by only one node at any point in time, but ownership can be transferred between nodes. The result is that each node has access to all cluster configuration data.
- Every node communicates with the other nodes in the cluster through one or more network adapters that attach nodes to networks.

- Every node in the cluster is aware of another system joining or leaving the cluster.
- Every node in the cluster is aware of the resources that are running on all nodes in the cluster.
- All nodes in the cluster are grouped under a common cluster name, which is used when accessing and managing the cluster.

Table 5-1 defines various states of a node that can occur in cluster operation.

Table 5-1. Node States and Definitions

State	Definition
Down	The node is not actively participating in cluster operations.
Joining	The node is in the process of becoming an active participant in the cluster operations.
Paused	The node is actively participating in cluster operations but cannot take ownership of resource groups and cannot bring resources online.
Up	The node is actively participating in all cluster operations, including hosting cluster groups.
Unknown	The state cannot be determined.

When the Cluster Service is installed for the first time on a node, the administrator must choose whether that node forms its own cluster or joins an existing cluster. When the Cluster Service is started on a node, that node searches for other active nodes on networks enabled for internal cluster communications.

Forming a New Cluster

If a node cannot join a cluster, the node attempts to form the cluster by gaining control of the quorum disk. If the node gains control of the quorum disk, the node forms the cluster and uses the recovery logs in the quorum disk to update its cluster database. The Cluster Service maintains a consistent, updated copy of the cluster database on all active nodes.

Joining an Existing Cluster

A node can join an existing cluster if it can communicate with another cluster node. If a cluster exists and the joining node finds an active node, it attempts to join that node's cluster. If it succeeds, the Cluster Service then validates the node's name and verifies version compatibility. If the validation process succeeds, the node joins the cluster. The node is updated with the latest copy of the cluster database.

Groups

A group is a collection of cluster resources with the following characteristics:

- All of the resources in the group are moved to the alternate node when one resource in a group fails and it is necessary to move the resource to an alternate node.
- A group is always owned by one node at any point in time, and a resource is always a member of a single group. Therefore, all of a group's resources reside on the same node.

Groups enable resources to be combined into larger logical units. Typically a group is made up of related or dependent resources, such as applications and their associated peripherals and data. However, groups can also be established with resources that are unrelated and nondependent to balance the load or for administrative convenience.

Every group maintains a prioritized list of the nodes that can and should act as its host. The preferred nodes list is generated by the Cluster Service. Cluster Service produces a list of preferred nodes for a group from the list of possible owners that is maintained by the group's resources and can be modified by an Administrator.

To maximize the processing power of a cluster, establish at least as many groups as there are nodes in the cluster.

Cluster Resources

A cluster resource is any physical or logical component that has the following characteristics:

- Can be brought online and taken offline
- Can be managed in a server cluster
- Can be hosted (owned) by only one node at a time

To manage resources, the Cluster Service communicates to a resource Dynamic Link Libraries (DLL) through a Resource Monitor. When the Cluster Service makes a request of a resource, the Resource Monitor calls the appropriate entry-point function in the resource DLL to check and control the resource's state.

Dependent Resources

A dependent resource requires—or depends on—another resource to operate. For example, if a Generic Application resource requires access to clustered physical storage, it would depend on a physical disk resource. A resource can specify one or more resources on which it is dependent; it can also specify a list of nodes on which it is able to run.

The following terms describe resources in a dependency relationship:

- **Dependent resource** — A resource that depends on other resources (the dependencies).
- **Dependency** — A resource on which another resource depends.
- **Dependency tree** — A series of dependency relationships such that resource A depends on resource B, resource B depends on resource C, and so on.

  Resources in a dependency tree obey the following rules:

  – A dependent resource and all of its dependencies must be in the same group.

  – The Cluster Service takes a dependent resource offline before any of its dependencies are taken offline and brings a dependent resource online after all its dependencies are online, as determined by the dependency hierarchy.

Setting Resource Properties

Using the resource **Properties** dialog box, you can perform the following tasks:

- View or change the resource name
- View or change the resource description and possible owners
- Assign a separate memory space for the resource
- View the resource type, group ownership, and resource state
- View which node currently owns the resource
- View pre-existing dependencies and modify resource dependencies
- Specify whether to restart a resource and the settings used to restart the resource (if required)
- Check the online state of the resource by configuring the **Looks Alive** and **Is Alive** polling intervals in the Cluster Service
- Specify the time requirement for resolving a resource in a pending state (**Online Pending** or **Offline Pending**) before the Cluster Service places the resource in **Offline** or **Failed** status
- Set specific resource parameters

The **General**, **Dependencies**, and **Advanced** tabs are the same for every resource. Some resource types support additional tabs.

Properties of a cluster object should not be updated on multiple nodes simultaneously. See the MSCS online documentation for more information.

Configuring Resource Dependencies

Groups function properly only if resource dependencies are configured correctly. The Cluster Service uses the dependencies list when bringing resources online and offline. For example, if a group in which a physical disk and a file share are located is brought online, the physical disk containing the file share must be brought online before the file share.

Table 5-2 shows resources and their dependencies. The resources in the right column must be configured before you create the resource.

Table 5-2. Cluster Resources and Required Dependencies

Resource	Required Dependencies
File share	Network name (only if configured as a distributed file system [DFS] root)
IP address	None
Network name	IP address that corresponds to the network name
Physical disk	None

Setting Advanced Resource Properties

You can configure the advanced resource properties using the **Advanced** tab in the resource **Properties** dialog box. Use the **Advanced** tab to have the Cluster Service perform the following tasks:

- Restart a resource or allow the resource to fail.
  - To restart the resource, select **Affect the group** (if applicable).
  - To fail over the resource group to another cluster node when the resource fails, select **Affect the group** and then enter the appropriate values in **Threshold** and **Period**. If you do not select **Affect the group**, the resource group will not fail over to the healthy cluster node.

  The **Threshold** value determines the number of attempts by the Cluster Service to restart the resource before the resource fails over to a healthy cluster node.

  The **Period** value assigns a time requirement for the **Threshold** value to restart the resource.
- Adjust the time parameters for **Looks Alive** (general check of the resource) or **Is Alive** (**detailed check of the resource**) to determine if the resource is in the online state.
- Select the default number for the resource type.

  To apply the default number, select **Use resource type value**.
- Specify the time parameter for a resource in a pending state (**Online Pending** or **Offline Pending**) to resolve its status before moving the resource to **Offline** or **Failed** status.

Resource Parameters

The **Parameters** tab in the **Properties** dialog box is available for most resources. Table 5-3 lists each resource and its configurable parameters.

Table 5-3. Resources and Configurable Parameters

Resource	Configurable Parameters
File share	Share permissions and number of simultaneous users
	Share name (clients will detect the name in their browse or explore lists)
	Share comment
	Shared file path
IP address	IP address
	Subnet mask
	Network parameters for the IP address resource (specify the correct cluster network)
Network name	System name
Physical disk	Drive for the physical disk resource (the drive cannot be changed after the resource is created)

Quorum Disk (Quorum Resource)

The quorum resource is a common resource in the cluster that is accessible by all of the cluster nodes. Normally a physical disk on the shared storage, the quorum resource maintains data integrity, cluster unity, and cluster operations—such as forming or joining a cluster—by performing the following tasks:

- **Enables a single node to gain and defend its physical control of the quorum resource** — When the cluster is formed or when the cluster nodes fail to communicate, the quorum resource guarantees that only one set of active, communicating nodes is allowed to form a cluster.
- **Maintains cluster unity** — The quorum resource allows cluster nodes that can communicate with the node containing the quorum resource to remain in the cluster. If a cluster node fails for any reason and the cluster node containing the quorum resource is unable to communicate with the remaining nodes in the cluster, MSCS automatically shuts down the node that does not control the quorum resource.
- **Stores the most current version of the cluster configuration database and state data** — If a cluster node fails, the configuration database helps the cluster recover a failed resource or recreates the cluster in its current configuration.

The only type of resource supported by the Cluster Service that can act as a quorum resource is the physical disk resource. However, developers can create their own quorum disk types for any resources that meet the arbitration and storage requirements.

Using the Quorum Disk for Cluster Integrity

The quorum disk is also used to ensure cluster integrity by performing the following functions:

- Maintaining the cluster node database
- Ensuring cluster unity

When a node joins or forms a cluster, the Cluster Service must update the node's private copy of the cluster database. When a node joins an existing cluster, the Cluster Service can retrieve the data from the other active nodes. However, when a node forms a cluster, no other node is available. The Cluster Service uses the quorum disk's recovery logs to update the node's cluster database, thereby maintaining the correct version of the cluster database and ensuring that the cluster is intact.

For example, if node 1 fails, node 2 continues to operate, writing changes to the cluster database. Before you can restart node 1, node 2 fails. When node 1 becomes active, it updates its private copy of the cluster database with the changes made by node 2 using the quorum disk's recovery logs to perform the update.

To ensure cluster unity, the operating system uses the quorum disk to ensure that only one set of active, communicating nodes is allowed to operate as a cluster. A node can form a cluster only if it can gain control of the quorum disk. A node can join a cluster or remain in an existing cluster only if it can communicate with the node that controls the quorum disk.

For example, if the private network (cluster interconnect) between cluster nodes 1 and 2 fails, each node assumes that the other node has failed, causing both nodes to continue operating as the cluster. If both nodes were allowed to operate as the cluster, the result would be two separate clusters using the same cluster name and competing for the same resources. To solve this problem, MSCS uses the node that owns the quorum disk to maintain cluster unity and solve this problem. In this scenario, the node that gains control of the quorum disk is allowed to form a cluster, and the other fails over its resources and becomes inactive.

Resource Failure

A failed resource is not operational on the current host node. At periodic intervals, the Cluster Service checks to see if the resource appears operational by periodically invoking the Resource Monitor. The Resource Monitor uses the resource DLL for each resource to detect if the resource is functioning properly. The resource DLL communicates the results back through the Resource Monitor to the Cluster Service.

Adjusting the Poll Intervals

You can specify how frequently the Cluster Service checks for failed resources by setting the **Looks Alive (general resource check)** and **Is Alive (detailed resource check)** poll intervals. The Cluster Service requests a more thorough check of the resource's state at each **Is Alive** interval than it does at each **Looks Alive** interval; therefore, the **Is Alive** poll interval is typically longer than the **Looks Alive** poll interval.

NOTE: Do not adjust the **Looks Alive** and **Is Alive** settings unless instructed by technical support.

Adjusting the Threshold and Period Values

If the resource DLL reports that the resource is not operational, the Cluster Service attempts to restart the resource. You can specify the number of times the Cluster Service can attempt to restart a resource in a given time interval. If the Cluster Service exceeds the maximum number of restart attempts (**Threshold** value) within the specified time period (**Period** value), and the resource is still not operational, the Cluster Service considers the resource to be failed.

NOTE: See "Setting Advanced Resource Properties" to configure the **Looks alive, Is alive, Threshold,** and **Period** values for a particular resource.

NOTE: Do not adjust the **Threshold** and **Period** values settings unless instructed by technical support.

Configuring Failover

You can configure a resource to fail over an entire group to another node when a resource in that group fails for any reason. If the failed resource is configured to cause the group that contains the resource to fail over to another node, Cluster Service will attempt a failover. If the number of failover attempts exceeds the group's threshold and the resource is still in a failed state, the Cluster Service will attempt to restart the resource. The restart attempt will be made after a period of time specified by the resource's **Retry Period On Failure** property, a property common to all resources.

When you configure the **Retry Period On Failure** properly, consider the following guidelines:

- Select a unit value of minutes, rather than milliseconds (the default value is milliseconds).
- Select a value that is greater or equal to the value of the resource's restart period property. This rule is enforced by the Cluster Service.

NOTE: Do not adjust the **Retry Period On Failure** settings unless instructed by technical support.

Creating a New Resource

Before you add a resource to your PowerEdge Cluster, you must verify that the following elements exist in your cluster:

- The type of resource is either one of the basic types provided with MSCS or a custom resource type provided by the application vendor, Microsoft, or a third party vendor.
- A group that contains the resource already exists within your cluster.
- All dependent resources have been created.
- A separate Resource Monitor—recommended for any resource that has caused problems in the past.

To create a new resource:

1. Click the **Start** button and select **Programs→ Administrative Tools→ Cluster Administrator**.

   The **Cluster Administrator** window appears.

2. In the console tree (usually the left pane), double-click the **Groups** folder.

3. In the details pane (usually the right pane), click the group to which you want the resource to belong.
4. On the **File** menu, point to **New**, and then click **Resource**.
5. In the New Resource wizard, type the appropriate information in **Name** and **Description**, and click the appropriate information in **Resource type** and **Group**.
6. Click **Next**.
7. Add or remove possible owners of the resource, and then click **Next**.

   The **New Resource** window appears with **Available resources** and **Resource dependencies** selections.
8. To *add* dependencies, under **Available resources**, click a resource, and then click **Add**.
9. To *remove* dependencies, under **Resource dependencies**, click a resource, and then click **Remove**.
10. Repeat step 7 for any other resource dependencies, and then click **Finish**.
11. Set the resource properties.

    For more information on setting resource properties, see the MSCS online help.

Deleting a Resource

1. Click the **Start** button and select **Programs→ Administrative Tools→ Cluster Administrator**.

   The **Cluster Administrator** window appears.
2. In the console tree (usually the left pane), click the **Resources** folder.
3. In the details pane (usually the right pane), click the resource you want to remove.
4. In the **File** menu, click **Delete**.

When you delete a resource, Cluster Administrator also deletes all the resources that have a dependency on the deleted resource.

File Share Resource Type

If you want to use a PowerEdge Cluster as a high-availability file server, you will need to select the type of file share for your resource. Three ways to use this resource type are available:

- **Basic file share** — Publishes a single file folder to the network under a single name.
- **Share subdirectories** — Publishes several network names—one for each file folder and all of its immediate subfolders. This method is an efficient way to create large numbers of related file shares on a single file server. For example, you can create a file share for each user with files on the cluster node.
- **DFS root** — Creates a resource that manages a stand-alone DFS root. Fault-tolerant DFS roots cannot be managed by this resource. A DFS root file share resource has required dependencies on a network name and an IP address. The network name can be either the cluster name or any other network name for a virtual server.

Configuring Active and Passive Cluster Nodes

Active nodes process application requests and provide client services. Passive nodes are backup nodes that ensure that client applications and services are available if a hardware or software failure occurs. Cluster configurations may include both active and passive nodes.

NOTE: Passive nodes must be configured with appropriate processing power and storage capacity to support the resources that are running on the active nodes.

Your cluster solution supports variations of active/active (active^x) and active/passive (active^x/passive^x) configurations. The variable *x* indicates the number of nodes that are active or passive.

Cluster solutions running the Windows operating system support active/active and active/passive configurations.

An active/active (active^x) configuration contains virtual servers running separate applications or services on each node. When an application is running on node 1, the remaining node(s) do not have to wait for node 1 to fail. Those node(s) can run their own cluster-aware applications (or another instance of the same application) while providing failover for the resources on node 1. For example, multiway failover is an active/active failover solution because running applications from a failed node can migrate to multiple active nodes in the cluster. However, you must ensure that adequate resources are available on each node to handle the increased load if one node fails.

- In an active/passive (active^x/passive^x) configuration, one or more *active* cluster nodes are processing requests for a clustered application while the *passive* cluster nodes only wait for the active node(s) to fail.

Failover and Failback

This section provides information about the failover and failback capabilities of the Cluster Service.

Failover

When an individual application or user resource (also known as a cluster resource) fails on a cluster node, the Cluster Service will detect the application failure and try to restart the application on the cluster node. If the restart attempt reaches a preset threshold, the Cluster Service brings the running application offline, moves the application and its resources to another cluster node, and restarts the application on the other cluster node(s). This process of automatically moving resources from a failed cluster node to other healthy cluster node(s) is called *failover*.

In order to fail over and fail back running applications, cluster resources are placed together in a group so the Cluster Service can move the cluster resources as a combined unit. This process ensures that the failover and/or failback procedures transfer all of the user resources as transparently as possible.

After failover, the Cluster Administrator can reset the following recovery policies:

- Application dependencies
- Application restart on the same cluster node
- Workload rebalancing (or failback) when a failed cluster node is repaired and brought back online

Failover Process

The Cluster Service attempts to fail over a group when any of the following conditions occur:

- The node currently hosting the group becomes inactive for any reason.
- One of the resources within the group fails, and it is configured to affect the group.
- Failover is forced by the System Administrator.

When a failover occurs, the Cluster Service attempts to perform the following procedures:

- The group's resources are taken offline.

  The resources in the group are taken offline by the Cluster Service in the order determined by the group's dependency hierarchy: dependent resources first, followed by the resources on which they depend.

  For example, if an application depends on a Physical Disk resource, the Cluster Service takes the application offline first, allowing the application to write changes to the disk before the disk is taken offline.

- The resource is taken offline.

  Cluster Service takes a resource offline by invoking, through the Resource Monitor, the resource DLL that manages the resource. If the resource does not shut down within a specified time limit, the Cluster Service forces the resource to shut down.

- The group is transferred to the next preferred host node.

  When all of the resources are offline, the Cluster Service attempts to transfer the group to the node that is listed next on the group's list of preferred host nodes.

  For example, if cluster node 1 fails, the Cluster Service moves the resources to the next cluster node number, which is cluster node 2.

- The group's resources are brought back online.

  If the Cluster Service successfully moves the group to another node, it tries to bring all of the group's resources online. Failover is complete when all of the group's resources are online on the new node.

The Cluster Service continues to try and fail over a group until it succeeds or until the number of attempts occurs within a predetermined time span. A group's failover policy specifies the maximum number of failover attempts that can occur in an interval of time. The Cluster Service will discontinue the failover process when it exceeds the number of attempts in the group's failover policy.

Modifying Your Failover Policy

Because a group's failover policy provides a framework for the failover process, make sure that your failover policy is appropriate for your particular needs. When you modify your failover policy, consider the following guidelines:

- Define the method in which the Cluster Service detects and responds to individual resource failures in a group.
- Establish dependency relationships between the cluster resources to control the order in which the Cluster Service takes resources offline.
- Specify **Time-out**, failover **Threshold**, and failover **Period** for your cluster resources.
  - **Time-out** controls how long the Cluster Service waits for the resource to shut down.
  - **Threshold** and **Period** control how many times the Cluster Service attempts to fail over a resource in a particular period of time.
- Specify a **Possible owner list** for your cluster resources. The **Possible owner list** for a resource controls which cluster nodes are allowed to host the resource.

Failback

When the System Administrator repairs and restarts the failed cluster node, the opposite process occurs. After the original cluster node has been restarted and rejoins the cluster, the Cluster Service will bring the running application and its resources offline, move them from the failover cluster node to the original cluster node, and then restart the application. This process of returning the resources back to their original cluster node is called failback.

You can configure failback to occur immediately at any given time, or not at all. However, be sure to configure the failback time during your offpeak hours to minimize the effect on users, as they may see a delay in service until the resources come back online.

6

Upgrading Your System to a Cluster Configuration

This section provides additional procedures for upgrading your non-clustered system to a Dell™ PowerEdge™ Cluster SE500W solution with the cluster components currently installed in your system. These cluster configurations include the Microsoft® Windows® 2000 Advanced Server, and Windows Server™ 2003 Enterprise Edition operating systems with MSCS.

Certification Requirements

All hardware components of the PowerEdge Cluster SE500W must pass Microsoft Windows Hardware Qualification Labs (MS WHQL) certification. Each cluster configuration must be certified and qualified against the latest MSCS certification program.

Before You Begin

Before you upgrade your non-clustered system to a Cluster SE500W solution:

- Back up the data on your cluster nodes.
- Verify that your cluster hardware and storage systems meet the minimum system requirements for a Cluster SE500W as described in "System Requirements."
- Verify that your cluster hardware and storage systems are installed and configured properly as explained in the following sections:
  - "Cabling Your Cluster Hardware"
  - "Preparing Your Systems for Clustering"

Dell certifies and supports only Cluster SE500W solutions that are configured with the Dell products described in this guide. For a description of the PowerEdge cluster components, see the *Platform Guide*.

Configurations Using Non-Dell Products

Configurations using non-Dell products, such as servers, rack cabinets, and storage systems, have not been approved by any safety agencies. It is the customer's responsibility to have these systems evaluated for any safety hazards or implications by a certified safety agency.

NOTE: Your system is safety-certified as a free-standing unit and as a component for use in a rack cabinet using the customer rack kit when both the rack cabinet and rack kit were designed for your system. The installation of your system and rack kit in any other rack cabinet has not been approved by any safety agencies. It is your responsibility to have the final combination of system and rack kit in a cabinet evaluated for suitability by a certified safety agency. The manufacturer disclaims all warranties and liability in connection with such combinations.

Completing the Upgrade

After you install the required hardware and network adapter upgrades, you can set up and cable the system hardware.

The final phase for upgrading to a Cluster SE500W solution is to install and configure Windows 2000 Advanced Server, or Windows Server 2003 with MSCS.

Upgrading Your Operating System

You can upgrade your Windows 2000 Advanced Server cluster to a Windows Server 2003 Enterprise Edition cluster using one of the following methods:

- **Standard upgrade** — Upgrading the operating system on each cluster node while all cluster nodes are disconnected from the client network. This procedure requires you to recreate your cluster configuration.

  See your Windows operating system documentation for performing a standard upgrade.

- **Rolling upgrade** — Upgrading the operating system on each cluster node while one cluster node is connected to the client network and available to handle client requests. This procedure *does not* require you to recreate your cluster configuration. However, each cluster node must be configured with the appropriate resources to run all virtual servers and services for the entire cluster while you upgrade the remaining node.

The following section explains how to perform a rolling upgrade on a two-node cluster running Windows 2000 Advanced Server.

NOTE: Only new (full) installations of the Microsoft Windows Server 2003 Enterprise x64 Edition are permitted. You cannot upgrade from existing operating system to the Windows Server 2003 Enterprise x64 Edition.

Performing a Rolling Upgrade

Before you perform a rolling upgrade:

- Ensure that your cluster nodes are running Windows 2000 Advanced Server.
- Backup your data and system states.
- Run the Check System Compatibility Wizard to determine if your cluster nodes are configured with the appropriate resources to run Windows Server 2003 Enterprise Edition.
- Ensure that your cluster service account user privileges are set to **Act as part of operating system**.

Upgrading Node 1

1. Click the **Start** button and select **Programs→ Administrative Tools→ Cluster Administrator**.
2. In **Cluster Administrator**, click one of the nodes in the cluster.

   The following steps refer to the node you selected as node 1.
3. Right-click a node and then click **Pause**.
4. Right-click a cluster group and then click **Move Group**.

   The Cluster Group is moved and restarted on node 2.
5. Repeat step 4 for the remaining cluster groups.
6. Insert the *Microsoft Windows Server 2003 Enterprise Edition* CD into the CD drive.
7. Double-click **Install Windows Server 2003 Enterprise Edition.**

   The **Windows Setup Wizard** window appears.
8. Follow the instructions in the Windows Setup Wizard to upgrade your operating system.

   NOTE: If you are running IIS World Wide Web Publishing service on your cluster node, this service is disabled during the upgrade to protect your system.
9. Verify that the upgraded node is functioning correctly.
   a. Click the **Start** button and select **Programs→ Administrative Tools→ Cluster Administrator**.
   b. Move one or more cluster resource groups from node 2.
   c. Verify that the resource group(s) can be brought online.
   d. Close **Cluster Administrator**.
10. Remove the *Enterprise Edition* CD from the CD drive.
11. Go to "Upgrading Node 2."

NOTE: After you upgrade node 1, your cluster is running two separate operating systems. Dell recommends that you do not modify your cluster configuration—such as adding or removing cluster nodes or resources—until you upgrade both cluster nodes.

Upgrading Node 2

1. On node 2, click the **Start** button and select **Programs→ Administrative Tools→ Cluster Administrator**.
2. In **Cluster Administrator**, right click node 1, and then click **Resume Node**.
3. Right-click node 2 and then click **Pause Node**.
4. Right-click a cluster group and then click **Move Group**.

   The cluster group is moved and restarted on node 1.
5. Repeat step 4 for the remaining cluster groups.
6. Insert the *Microsoft Windows Server 2003 Enterprise Edition* CD into the CD drive.
7. Double-click **Install Windows Server 2003 Enterprise Edition**.

   The **Windows Setup Wizard** window appears.
8. Follow the instructions in the **Windows Setup Wizard** window to upgrade your operating system.

   NOTE: If you are running IIS World Wide Web Publishing service on your cluster node, this service is disabled during the upgrade to protect your system.
9. Verify that the upgraded node is functioning correctly.
   a. Click the **Start** button and select **Programs→ Administrative Tools→ Cluster Administrator**.
   b. Move one or more cluster resource groups from node 1.
   c. Verify that the resource group(s) can be brought online.
   d. Close **Cluster Administrator**.
10. In **Cluster Administrator**, redistribute the cluster groups to the appropriate cluster nodes.

7

Maintaining Your Cluster

This section provides the following maintenance procedures for systems running the Microsoft® Windows® 2000 Advanced Server, and Windows Server™ 2003 operating systems:

- Adding a network adapter to a cluster node
- Changing the IP address of a cluster node
- Uninstalling MSCS
- Running **chkdsk /f** on a quorum disk
- Recovering from a corrupt quorum disk
- Replacing a cluster node

Adding a Network Adapter to a Cluster Node

This procedure assumes that Windows 2000 Advanced Server, or Windows Server 2003 with the latest Windows Service Pack, and MSCS are installed on both cluster nodes.

NOTE: The IP addresses used in the following sections are examples only and do not represent of actual addresses to use. The IP addresses are 192.168.1.101 for the network adapter in the first node and 192.168.1.102 for the network adapter in the second node. The subnet mask for both nodes is 255.255.255.0.

1. Move all cluster resources from the cluster node you are upgrading to another node in the cluster.

   See the Cluster Service documentation for information about moving cluster resources to a specific node.

2. Shut down the cluster node you are upgrading and install the additional network adapters in that system.

   See the system *Installation and Troubleshooting Guide* for instructions about installing expansion cards in your system.

3 Boot to the Windows operating system.

Windows Plug and Play detects the new network adapter and installs the appropriate drivers.

NOTE: If Plug and Play *does not* detect the new network adapter, the network adapter is not supported.

a Update the network adapter drivers (if required).

b After the drivers are installed, click the **Start** button, select **Control Panel**, and then double-click **Network Connections**.

c In the **Connections** box, locate the new network adapter that you installed in the system.

d Right-click the new network adapter, and then select **Properties**.

e Assign a unique static IP address, subnet mask, and gateway.

4 Ensure that the network ID portion of the new network adapter's IP address is different from the other adapter.

For example, if the first network adapter in the node had an address of 192.168.1.101 with a subnet mask of 255.255.255.0, you might enter the following IP address and subnet mask for the second network adapter:

IP address: `192.168.2.102`

Subnet mask: `255.255.255.0`

5 Click **OK** and exit network adapter properties.

6 On the Windows desktop, click the **Start** button and select **Programs→ Administrative Tools→ Cluster Administrator**.

7 Click the **Network** tab.

8 Verify that a new resource called "New Cluster Network" appears in the window.

To rename the new resource, right-click the resource and enter a new name.

9 Move all cluster resources to another cluster node.

10 Repeat step 2 through step 9 on each cluster node.

NOTE: Ensure that you assign the new network adapter the same IP address as the second network adapter on the first node (for example, 192.168.2.101) as you did with the second node.

If the installation and IP address assignments have been performed correctly, all of the new network adapter resources appear online and respond successfully to **ping** commands.

Changing the IP Address of a Cluster Node on the Same IP Subnet

NOTE: If you are migrating your cluster nodes to a different subnet, take all cluster resources offline and then migrate all nodes together to the new subnet.

1. Open **Cluster Administrator**.
2. Stop MSCS on the cluster node.

   The Cluster Administrator utility running on the second cluster node indicates that the first node is down by displaying a red icon in the **Cluster Service** window.
3. Reassign the IP address.
4. If you are running DNS, verify that the DNS entries are correct (if required).
5. Restart MSCS on the cluster node.

   The cluster nodes re-establish their connection and the Cluster Administrator changes the node icon back to blue to show that the node is back online.

Uninstalling MSCS From Clusters Running Windows 2000 Advanced Server

1. Take all resource groups offline or move them to another cluster node.
2. Stop Cluster Service on the node that you want to uninstall.
3. Click the **Start** button and select **Settings→ Control Panel→ Add/Remove Programs**.
4. Select **Add/Remove Windows Components**.
5. Deselect the check box for Cluster Service and click **Next**.
6. Click **Finish**.
7. From the remaining node, click the **Start** button and select **Programs→ Administrative Tools→ Cluster Administrator**.
8. Right-click the node icon and select **Options→ Evict Node**.
9. Close **Cluster Administrator**.
10. Restart the Cluster Service.

Removing Nodes From Clusters Running Windows Server 2003

1. Take all resource groups offline or move them to another cluster node.
2. Click the **Start** button, select **Programs→ Administrative Tools,** and then double-click **Cluster Administrator**.
3. In **Cluster Administrator**, right-click the icon of the node you want to uninstall and then select **Stop Cluster Service**.
4. In **Cluster Administrator**, right-click the icon of the node you want to uninstall and then select **Evict Node**.

   If you cannot evict the node and the node is the last node in the cluster:

   a. Open a command prompt.

   b. Type the following:

   ```
   cluster node <node_name> /force
   ```

   where *<node_name>* is the cluster node you are evicting from the cluster.
5. Close **Cluster Administrator**.

Running chkdsk /f on a Quorum Disk

You cannot run the **chkdsk** command with the **/f** (fix) option on a device that has an open file handle active. Because MSCS maintains an open handle on the quorum resource, you cannot run **chkdsk /f** on the hard drive that contains the quorum resource.

To run **chkdsk /f** on a quorum resource's hard drive:

1. Move the quorum resource temporarily to another drive:
   a. Right-click the cluster name and select **Properties**.
   b. Click the **Quorum** tab.
   c. Select another disk as the quorum disk and press <Enter>.
2. Run **chkdsk /f** on the drive that previously stored the quorum resource.
3. Move the quorum disk back to the original drive.

Recovering From a Corrupt Quorum Disk

The quorum disk maintains the configuration data necessary for cluster recovery when a cluster node fails. If the quorum disk resource is unable to come online, the cluster will not start and all of the shared drives will be unavailable. If this situation occurs, and you need to run **chkdsk** on the quorum disk, you can start the cluster manually from the command line.

To start the cluster manually from a command prompt:

1. Open a command prompt window.
2. Select the cluster folder directory by typing one of the following:

   `cd \2000\cluster` (for Windows 2000 Advanced Server), or
   `cd \windows\cluster` (for Windows Server 2003)
3. Start the cluster in manual mode (on one node only) with no quorum logging by typing the following:

   ```
   Clussvc -debug -noquorumlogging
   ```

   MSCS starts.
4. Run **chkdsk /f** on the disk designated as the quorum resource.

   To run the **chkdsk /f** utility:

   a. Open a second command prompt window.

   b. Type:

   ```
   chkdsk /f
   ```
5. After the **chkdsk** utility completes, stop MSCS by pressing <Ctrl><c>.
6. Restart the cluster service.

   To restart MSCS from the **Services** console:

   a. Click the **Start** button and select **Programs→ Administrative Tools→ Services**.

   b. In the **Services** window, right-click **Cluster Service**.

   c. In the drop-down menu, click the **Start** button.

   To restart MSCS from the command prompt:

   a. Open the second command prompt window that you opened in step 4a.

   b. Type the following:

   ```
   Net Start Clussvc
   ```

   The Cluster Service restarts.

See the Microsoft Knowledge Base article 258078 located at the Microsoft Support website at **www.microsoft.com** for more information on recovering from a corrupt quorum disk.

Replacing a Cluster-Enabled Dell PERC RAID Adapter

If you have to replace a cluster-enabled PERC RAID adapter, perform the following steps:

1. Turn off the failed node.
2. Disconnect the failed PERC RAID adapter's cable from the shared storage system.

NOTICE: If you replace your PERC RAID adapter, ensure that you enable cluster mode on the replacement PERC RAID adapter before you connect the SCSI cables to the shared storage system.

NOTICE: A failed PERC 4/DC card must be replaced with a PERC4/DC card. A failed PERC 4e/DC card must be replaced with a PERC 4e/DC card.

3. Replace the failed PERC RAID adapter in the system without reconnecting the cable.
4. Turn on the system with the replaced PERC RAID adapter and run the BIOS configuration utility.

NOTICE: If you replace a PERC RAID adapter that will be connected to shared storage system, you must set the appropriate SCSI ID before you connect the SCSI cables to the shared storage system.

5. Change the SCSI ID so that it differs from the SCSI ID on the peer cluster node's PERC RAID adapter.

   NOTE: See your PERC RAID adapter documentation for more information about changing the SCSI ID. Also, see the cluster configuration tables (if you completed the information in the tables) in the *Dell PowerEdge Cluster SE500W Platform Guide*.

6. Shut down the system.
7. Reconnect the system to the shared storage system.
8. Restart the system and restore the RAID configuration using configuration information stored on the disks. See the PERC RAID adapter documentation for more information about this procedure.

Replacing a Cluster Node

This section assumes that you have a recent tape backup of the cluster node that contains the local registry information.

1. Ensure that the replacement cluster node is physically disconnected from the storage system.
2. Ensure that Windows and the latest Windows Service Pack are installed and configured properly on the replacement node.
3. Install the correct network adapter drivers, assign the appropriate IP addresses, and install the PERC RAID adapter driver.
4. Shut down the replacement node.
5. Connect the SCSI cables from each PERC RAID adapter to the Dell™ PowerVault™ 22*x*S storage system.

6 Turn on the replacement cluster node.

If the PERC RAID adapter has been replaced, the following error message is reported:

```
Configuration of NVRAM and drives mismatch (Normal mismatch)
Run View/Add Configuration option of Config Utility
Press <Ctrl><H> for WebBIOS
Press A Key to Run Configuration Utility
Or <Alt><F10> to Continue
```

7 Press any key to enter the RAID controller's BIOS configuration utility, and select **Configure→ View/Add Configuration→ View Disk Configuration**.

Verify that the configuration that displays includes the existing configuration on the disks.

8 Press <Esc>, and select **Yes** to save the disk configuration, and exit the configuration utility.

9 Restart the system and allow Windows to start normally.

10 Add the replacement node to the cluster.

11 Use Cluster Administrator to verify that the node rejoins the cluster, and check the **Windows Event Viewer** to ensure errors were not encountered.

12 Reinstall any cluster applications (such as Microsoft SQL Server or Exchange Server onto the new node, if required).

NOTE: You may need to reinstall or configure your cluster applications before moving or testing the failover capabilities of any cluster resources to the new node.

Changing the Cluster Service Account Password in Windows Server 2003

To change the cluster service account password for all nodes in a cluster running Windows Server 2003, open a command prompt and type the following syntax:

```
Cluster /cluster:[cluster_name] /changepass
```

where `cluster_name` is the name of your cluster.

For help with changing the cluster password, type the following:

```
cluster /changepass /help
```

NOTE: Blank passwords are not allowed as cluster service account passwords in Windows Server 2003.

Reformatting a Cluster Disk

NOTE: Ensure that all client systems are disconnected from the cluster disk before you perform this procedure.

1. Click the **Start** button and select **Programs**→ **Administrative Tools**→ **Cluster Administrator**.
2. In the **Cluster Administrator** left window pane, expand the **Groups** directory.
3. In the **Groups** directory, right-click a cluster resource group that contains the disk to be reformatted, and select **Take Offline**.
4. In the **Cluster Administrator** right window pane, right-click the physical disk you are reformatting and select **Bring Online**.
5. In the **Cluster Administer** right window pane, right-click the physical disk you are reformatting and select **Properties**.

   The **Properties** window appears.
6. Click the **Advanced** tab.
7. In the **Advanced** tab menu in the **"Looks Alive" poll interval** box, select **Specify value**.
8. In the **Specify value** field, type:

   `6000000`

   where `6000000` equals 6,000,000 milliseconds (or 100 minutes).
9. Click **Apply**.
10. On the Windows desktop, right-click **My Computer** and select **Manage**.

    The **Computer Management** window appears.
11. In the **Computer Management** left window pane, click **Disk Management**.

    The physical disk information appears in the right window pane.
12. Right-click the disk you want to reformat and select **Format**.

    Disk Management reformats the disk.
13. In the **File** menu, select **Exit**.
14. In the **"Looks Alive" poll interval** box, select **Use value from resource type** and click **OK**.
15. In the **Cluster Administrator** left window pane, right-click the cluster group that contains the reformatted disk and select **Bring Online**.
16. In the **File** menu, select **Exit**.

Adding New Physical Drives to an Existing Shared Storage System

The Dell™ PowerEdge™ Cluster SE500W solutions consist of two systems that share an external SCSI storage enclosure PowerVault 22*x*S storage system. Each system contains a PERC RAID adapter with cluster-enabled firmware. The following procedure describes adding additional storage to an existing shared storage system in the cluster configuration.

To add new physical drives to an existing shared storage system in the cluster, perform the following steps:

1. Stop all I/O activity.
2. Ensure that both nodes are online.
3. Install the new physical hard drives into the PowerVault 22*x*S storage system.

CAUTION: See the *Dell PowerVault 220S and 221S Installation and Troubleshooting Guide*, which provides safety instructions for installing components into the PowerVault 22*x*S storage system.

4. Restart node 1 and press <Ctrl><m> during the system POST to launch the PERC RAID adapter BIOS Configuration utility.
5. Configure the virtual disks.

   NOTE: See the PERC RAID adapter documentation.

6. Restart node 1.
7. After system restarts, use Disk Manager to write the disk signature, create a new partition, assign drive letters, and format the partition with NTFS.
8. Restart node 1.
9. On node 1, use Cluster Administrator to add a new group (for example Disk Group *n*:).
10. Select possible owners, but do not bring the group online yet.
11. Add a new resource (for example, Disk *z*:).
12. Select **Physical Disk** for the type of resource, and assign it to the new group you just created.
13. Select possible owners, and select the drive letter that you assigned to the new array.
14. Bring the new group that you just added online.
15. Reboot node 2, and ensure that node 2 is completely online before you continue.
16. To verify that the new resource group is online and the drive is accessible using the cluster name, connect to **clustername***n*$, where *n* is the drive letter you assigned to the newly added disk, and use Cluster Administrator to verify that you can move the new disk group to the other cluster node.

Rebuilding Operation in Dell OpenManage Utilities

For the rebuild operation, see your Dell OpenManage™ Array Manager or Dell OMSM documentation.

If the cluster node is rebooted or power to the node is lost while a PERC RAID adapter is rebuilding a shared array, the controller terminates the rebuild operation and identifies the hard drive as failed. This condition also occurs if the rebuild is performed from the PERC RAID adapter BIOS Configuration utility and the user exits the utility before the rebuild completes. This condition occurs with all versions of the PERC RAID adapter firmware on both standard and cluster-enabled controllers.

If the second node in the clustered configuration is turned on, it restarts the operation.

If the rebuild fails to complete due to a system restart, the rebuild must be reinitiated using the PERC RAID adapter BIOS configuration utility or using the Array Manager running under the appropriate Microsoft Windows operating system.

Upgrading the PowerVault 22*x*S EMM Firmware Using Array Manager

NOTE: Before upgrading the EMM firmware, ensure that all I/O activity is suspended, and shut down the node you will not be working from when you upgrade the EMM firmware. If you do not shut down the second node, the firmware on the EMM attached to that node may not be updated.

NOTE: Use Array Manager version 3.7 or later to perform the following procedure.

To download the PowerVault 22*x*S EMM firmware onto a cluster node:

1. Download the latest EMM firmware from the Dell Support website (located at **support.dell.com**) to your hard drive or to a diskette.
2. Shut down node B.
3. Stop all I/O activity on node A.
4. Launch the Array Manager Console from node A by clicking the **Start** button and selecting **Programs→ Dell OpenManage Applications→ Array Manager→ Array Manager console.**
5. In the **Arrays** directory, select **PERC Subsystem 1→ PERC RAID Adapter Controller** *x* **<Cluster>→ <Channel 0>** or **<Channel 1>**.

   where *x* indicates the number associated with the controller on the system. Select the channel (0 or 1) to which the enclosure is attached.
6. If you downloaded the EMM firmware to a diskette, ensure that the diskette is inserted.
7. Right-click the enclosure icon for the desired channel, and select **Download Firmware**.

   You can also click the channel number and select **Download Firmware** from the **Task Menu**.
8. From the **Firmware Download** dialog box, click **Browse** and navigate to the EMM firmware that you downloaded to your hard drive or diskette.

9 Verify that the selected file is correct.

10 Click **Download Firmware** to begin the download process.

This process takes several minutes to complete.

11 When the message `Firmware Downloaded Successfully` appears, click **OK**.

12 Repeat steps 3 through 9 for each channel that has an enclosure attached.

13 To verify the firmware upgrade for each channel, right-click the channel number, select **Properties**, and view the version information.

14 Start up node B and resume I/O activity.

A

Troubleshooting

This appendix provides troubleshooting information for Dell™ PowerEdge™ Cluster SE500W configurations.

Table A-1 describes general cluster problems you may encounter and the probable causes and solutions for each problem.

Table A-1. General Cluster Troubleshooting

Problem	Probable Cause	Corrective Action
The RAID drives in the Dell PowerVault™ 22*x*S storage system are not accessible by one of the cluster nodes, or the shared storage system is not functioning properly with the cluster software.	The SCSI cables are loose or defective, or the cables exceed the maximum allowable length.	Check the cable connections or replace the cable with a working cable. For more information on the length of SCSI cables, see "Cabling Your Cluster Hardware."
A disk resource will not move over to another node or will not come online.	The PERC RAID adapters connected to a single storage system are not configured consistently.	Ensure that the RAID configuration is identical for each channel between the PERC 3/DC cards connected to a shared storage system. Ensure that cluster mode is enabled on both PERC RAID adapters and that their SCSI IDs are different on each node.
	The PowerVault 22*x*S is not running in cluster mode.	Configure the PowerVault 22*x*S for cluster mode. For more information, see "Preparing Your Systems for Clustering."
	If the cluster has multiple PowerVault 22*x*S storage systems, the cabling between the PERC RAID adapter and the storage systems is wrong.	Ensure that the cables attached to each channel of the PERC RAID adapter in each server node are connected to the correct storage system and that the channels on an optional second PERC RAID adapter in each server node are connected to the correct system.
		Attach or replace the SCSI cable between the cluster node and the shared storage system.

Table A-1. General Cluster Troubleshooting *(continued)*

Problem	Probable Cause	Corrective Action
	Enclosure management modules (EMMs) are not installed.	Install EMMs.
	The PERC RAID adapter drivers are not installed in your Microsoft® Windows® operating system.	Install the drivers. See the PERC RAID adapter documentation for more information.
The option to change the SCSI IDs is not visible in the PERC 3/DC BIOS.	Cluster mode is not enabled.	Enabling cluster mode will permit you to change the SCSI IDs.
One or more of the SCSI controllers are not detected by the system. PERC RAID adapters hang during boot.	The controllers for the shared storage system have the same SCSI ID as their peer adapters in the other system (that is, the same SCSI ID as the controllers connected to the other side of the shared storage system).	Change one of the controller SCSI IDs so that the ID numbers do not conflict. Set the controller in the primary node to SCSI ID 7 (default), and set the controller in the secondary node to SCSI ID 6. See the PERC RAID adapter documentation for more information about setting SCSI host adapter IDs.
Array Manager and PERC RAID adapter BIOS utility only report 13 drives in cluster mode.	Normal	The SCSI ID limitations are imposed by SCSI protocol. As a result of this limitation, the last slot in the PowerVault 22*x*S storage system cannot be utilized in cluster mode.
One of the nodes takes a long time to join the cluster, or fails to join the cluster.	The node-to-node network has failed due to a cabling or hardware failure.	Check the network cabling. Ensure that the node-to-node interconnection and the public network are connected to the correct network adapters.
	Long delays in node-to-node communications may be normal.	Verify that the nodes can communicate with each other by running the **ping** command from each node to the other node. Try both the host name and IP address when using the **ping** command.
	One or more nodes may have the Internet Connection Firewall enabled, and the RPC communication between the nodes has been inadvertently blocked.	Configure the Internet Connection Firewall to allow communications that are required by the cluster service and the clustered applications or services. For more information see KB883398 at **support.microsoft.com**.

Table A-1. General Cluster Troubleshooting *(continued)*

Problem	Probable Cause	Corrective Action
You are prompted to configure one network instead of two during MSCS installation.	The TCP/IP configuration is incorrect.	The node-to-node network and public network must be assigned static IP addresses on different subnets. See "Assigning Static IP Addresses to Your Cluster Resources and Components" for information about assigning the network IPs.
	The private (point-to-point) network is disconnected.	Ensure that both systems are powered on so that both network adapters in the private network are available.
Client systems are dropping off of the network while the cluster is failing over.	With MSCS, the service provided by the recovery group becomes temporarily unavailable to clients during failover. Clients may lose their connection if their attempts to reconnect to the cluster are too infrequent or if they end too abruptly.	The time that the service is temporarily unavailable varies depending on the application. Contact the application program vendor for more information.
Only one network segment appears during Cluster Service installation.	Public and private networks segments are not unique.	Place all installed network adapters in a cluster node on separate IP networks. Ensure that the same network segments that were used for each network adapter are identical on the second cluster node.

Table A-1. General Cluster Troubleshooting *(continued)*

Problem	Probable Cause	Corrective Action

Attempts to connect to a cluster using Cluster Administrator fail.	The Cluster Service has not been started.
	A cluster has not been formed on the system.
	The system has just been booted and services are still starting.
		Verify that the Cluster Service is running and that a cluster has been formed. Use the **Event Viewer** and look for the following events logged by the Cluster Service:
		`Microsoft Cluster Service successfully formed a cluster on this node.`
		or
		`Microsoft Cluster Service successfully joined the cluster.`
		If these events do not appear in Event Viewer, see the *Microsoft Cluster Service Administrator's Guide* for instructions on setting up the cluster on your system and starting the Cluster Service.

	The cluster network name is not responding on the network because the Internet Connection Firewall is enabled on one or more nodes.	Configure the Internet Connection Firewall to allow communications that are required by the cluster service and the clustered applications or services. For more information see KB883398 at **support.microsoft.com**.
Using Microsoft Windows NT® 4.0 to remotely administer a Windows 2000 Advanced Server or Windows Server™ 2003 cluster generates error messages.	Normal. Some resources in Windows 2000, Advanced Server and Windows Server 2003 are not supported in Windows NT 4.0.	Dell strongly recommends that you use Windows 2000 Professional, Windows 2000 Server, or Windows 2000 Advanced Server for remote administration of a cluster running Windows 2000 Advanced Server. Similarly, you should use Windows XP Professional or Windows Server 2003 for remote administration of a cluster running Windows Server 2003.
MSCS does not show any available shared disks during installation.	The PERC RAID adapter drivers are not installed in the operating system. Disks are configured as dynamic disks.	Install the drivers. See the PERC RAID adapter documentation for more information. Change disks to "basic" before cluster installation. See "Maintaining Your Cluster" for more information on configuring dynamic disks as basic disks.

Table A-1. General Cluster Troubleshooting *(continued)*

Problem	Probable Cause	Corrective Action
One of the nodes can access one of the shared hard drives, but the second node cannot.	If MSCS is installed, this situation is normal.	If MSCS is installed, only the node that owns the disk resource will be able to access the disk. The other node will show the disk resource as offline in Windows Disk Management.

NOTE: For support and troubleshooting information for Dell PowerVault™ 220S and 221S Systems refer to **support.dell.com/support/edocs/stor-sys/spv22xs/**. For support and troubleshooting information for PERC cards refer to **support.dell.com/support/edocs/storage/RAID/**.

Abbreviations and Acronyms

A
ampere(s)

API
Application Programming Interface

AC
alternating current

ACM
advanced cooling module

BBS
Bulletin Board Service

BDC
backup domain controller

BIOS
basic input/output system

bps
bits per second

BTU
British thermal unit

C
Celsius

cm
centimeter(s)

DC
direct current

DFS
distributed file system

DHCP
dynamic host configuration protocol

DLL
dynamic link library

DNS
domain naming system

ESD
electrostatic discharge

EMM
enclosure management module

ERP
enterprise resource planning

F
Fahrenheit

FC
Fibre Channel

FCAL
Fibre Channel arbitrated loop

ft
feet

FTP
file transfer protocol

g
gram(s)

GB
gigabyte

Gb
gigabit

Gb/s
gigabits per second

GUI
graphical user interface

HBA
host bus adapter

HSSDC
high-speed serial data connector

HVD
high-voltage differential

Hz
hertz

ID
identification

IIS
Internet Information Server

I/O
input/output

IP
Internet Protocol

Kb
kilobit(s)

KB
kilobyte(s)

KVM
Keyboard Video Mouse

lb
pound(s)

LAN
local area network

LED
light-emitting diode

LS
loop resiliency circuit/SCSI enclosure services

LVD
low-voltage differential

m
meter

MB
megabyte(s)

MB/sec
megabyte(s) per second

MHz
megahertz

MNS
Majority Node Set

MS WHQL
Microsoft Windows Hardware Qualification Labs

MSCS
Microsoft® Cluster Service

MSDTC
Microsoft Distributed Transaction Coordinator

NIC
Network Interface card

NLB
Network Load Balancing

NTFS
NT File System

NVRAM
nonvolatile random-access memory

OMSM
OpenManage Enhanced Storage Management

PAE
physical address extension

PCB
printed circuit board

PDC
primary domain controller

PDU
power distribution unit

PERC
PowerEdge™ Expandable RAID Controller

PERC 4/DC
PERC fourth generation, dual channel

PERC 4e/DC
PERC fourth generation express, dual channel

PCI
Peripheral Component Interconnect

POST
power-on self-test

RAID
redundant array of independent disks

RAM
random access memory

rpm
revolutions per minute

SAFTE
SCSI accessed fault-tolerant enclosures

SCSI
small computer system interface

sec
second(s)

SES
SCSI enclosure services

SMP
symmetric multiprocess

SNMP
Simple Network Management Protocol

TCP/IP
Transmission Control Protocol/Internet Protocol

UHDCI
ultra high-density connector interface

UPS
uninterruptible power supply

V
volt(s)

VHDCI
very high-density connector interface

WINS
Windows Internet Naming Service

B

Cluster Data Form

The configuration matrix and data form on the following pages are provided for the system installer to record pertinent information about Dell™ PowerEdge™ Cluster SE500W configurations. The data form is used for installing Microsoft® Windows® 2000 Advanced Server, and Windows Server™ 2003 operating systems on PowerEdge clusters.

Make a copy of the appropriate section of the data form to use for the installation or upgrade, complete the requested information on the form, and have the completed form available if you need to call Dell for technical assistance. If you have more than one cluster, complete a copy of the form for each cluster.

PowerEdge Cluster SE500W Solution Data Form

You can attach the following form to the back of each cluster node or rack. The system installer may want to use the form to record important information about the hardware on each cluster component. Have a copy of the form available any time you call Dell for technical support.

Cluster Type	PowerEdge Cluster SE500W Solution
Cluster name	
Domain name	
Cluster IP address	
Cluster subnet mask (same as public network)	
Cluster Service account	
Cluster Service password	
Installer	
Date installed	
Applications	
Location	
Notes	

Node (Server Name)	Server Type	Cluster Name	Service Tag Number
Node 1			
Node 2			

Network Settings	TCP/IP Address	Subnet Mask	Private or Public?
Node 1, network adapter 1			
Node 1, network adapter 2			
Additional Node 1 network adapter(s)			
Node 2, network adapter 1			
Node 2, network adapter 2			
Additional Node 2 network adapter(s)			

System	Storage 1	Storage 2	Storage 3	Storage 4	SCSI ID
Node 1, PERC RAID adapter 1					
Node 2, PERC RAID adapter 1					
Node 1, PERC RAID adapter 2					
Node 2, PERC RAID adapter 2					

Dell PowerVault™ 22*x*S Storage System	Description of Installed Items (Drive letters, RAID types, applications/data)
Storage 1	
Storage 2	
Storage 3	
Storage 4	

Component	Storage 1	Storage 2	Storage 3	Storage 4
Service Tag				

PCI Slot Number	Adapter Installed (PERC RAID adapter, network adapter, and so on)	Use (public network, private network, shared storage, internal drives)	PCI Slot Description
PCI slot 1			
PCI slot 2			
PCI slot 3			
PCI slot 4			
PCI slot 5			
PCI slot 6			
PCI slot 7			
PCI slot 8			
PCI slot 9			
PCI slot 10			
PCI slot 11			

Index

U

V

W

Figures

Tables